arXiv:1701.08306v1 [cs.AI] 28 Jan 2017

Practical Reasoning with Norms for Autonomous Software Agents (Full Edition)

Zohreh Shams[a,1,*], Marina De Vos[a], Julian Padget[a], Wamberto W. Vasconcelos[b]

[a]*Department of Computer Science, University of Bath, UK (BA2 6AH)*
[b]*Department of Computing Science, University of Aberdeen, UK (AB24 3UE)*

Abstract

Autonomous software agents operating in dynamic environments need to constantly reason about actions in pursuit of their goals, while taking into consideration norms which might be imposed on those actions. Normative practical reasoning supports agents making decisions about what is best for them to (not) do in a given situation. What makes practical reasoning challenging is the interplay between goals that agents are pursuing and the norms that the agents are trying to uphold. We offer a formalisation to allow agents to plan for multiple goals and norms in the presence of *durative* actions that can be executed *concurrently*. We compare plans based on decision-theoretic notions (i.e. utility) such that the utility gain of goals and utility loss of norm violations are the basis for this comparison. The set of *optimal* plans consists of plans that maximise the overall utility, each of which can be chosen by the agent to execute. We provide an implementation of our proposal in Answer Set Programming, thus allowing us to state the original problem in terms of a logic program that can be queried for solutions with specific properties. The implementation is proven to be sound and complete.

Keywords: Intelligent Agents, Practical Reasoning, Norms, Goals

*Corresponding author.
Email addresses: z.shams@bath.ac.uk (Zohreh Shams), m.d.vos@bath.ac.uk (Marina De Vos), j.a.padget@bath.ac.uk (Julian Padget), wvasconcelos@acm.org (Wamberto W. Vasconcelos)
[1]Present Address: Computer Laboratory, University of Cambridge, UK (CB3 0FD)

1. Introduction

Reasoning about what to do – known as practical reasoning – for an agent pursuing different goals is a complicated task. When conducting practical reasoning, the agents might exhibit undesirable behaviour that was not predicted. The necessity of controlling undesirable behaviour has given rise to the concept of norms that offer a way to define ideal behaviour for autonomous software agents in open environments. Such norms often define obligations and prohibitions that express what the agent is obliged to do and what the agent is prohibited from doing.

Depending on their computational interpretation, norms can be regarded as *soft* or *hard constraints*. When modelled as hard constraints, the agent subject to the norms is said to be *regimented*, in which case the agent has no choice but blindly to follow the norms (Esteva et al., 2001). Although regimentation guarantees norm compliance, it greatly restricts agent autonomy. Moreover, having individual goals to pursue, self-interested agents might not want to or might not be able to comply with the norms imposed on them. Conversely, *enforcement* approaches, in which norms are modelled as soft constraints, leave the choice of complying with or violating the norms to the agent. However, in order to encourage norm compliance, there are consequences associated, namely a punishment when agents violate a norm (López y López et al., 2005; Pitt et al., 2013) or a reward when agents comply with a norm (Aldewereld et al., 2006). In some approaches (*e.g.*, (Aldewereld et al., 2006; Alrawagfeh and Meneguzzi, 2014; Oren et al., 2011)) there is a utility gain/loss associated with respecting norm or not, whereas in the *pressured norm compliance* approaches (*e.g.*, (López y López et al., 2005)), the choice to violate a norm or not is determined by how the norm affects the satisfaction or hindrance of the agent's goals.

Existing work (e.g. (Oren et al., 2011; Panagiotidi et al., 2012a; Criado et al., 2010; Meneguzzi et al., 2015)) on normative practical reasoning using enforcement either consider plan generation or plan selection where there is a set of pre-generated plans available to the agent. In these works, the attitude agents have toward norms is often one of compliance, meaning that their plans are often selected or, in some approaches, customised, to ensure norm compliance (e.g., Kollingbaum (2005); Alechina et al. (2012); Oren et al. (2011)). We argue that in certain situations, an agent might be better off violating a norm which, if followed, would make it impossible for the agent to achieve an important goal or complying with a more important norm.

In this paper we set out an approach for practical reasoning that considers norms in both plan generation and plan selection. We extend current work on normative plan generation such that the agent attempts to satisfy a set of potentially conflicting goals in the presence of norms, as opposed to conventional planning problems that generate plans for a single goal (Oren et al., 2011; Panagiotidi et al., 2012a). Such an extension is made on top of STRIPS (Fikes and Nilsson, 1971), the most established planning domain language that lays the foundation of many automated planning languages. Additionally, since in reality the actions are often non-atomic, our model allows for planning with durative actions that can be executed concurrently. Through our practical reasoning process agents consider *all* plans (i.e., sequences of actions), including those leading to norm compliance and violation; each plan gets an associated overall utility for its sequence of actions, goals satisfied, and norms followed/violated, and agents can decide which of them to pursue by comparing the relative importance of goals and norms via their utilities. The plan an agent chooses to follow is not necessarily norm-compliant; however, our mechanism guarantees that the decision will maximise the overall plan utility, and this justifies the occasional violation of norms as the plan is followed. Both plan generation and plan selection mechanisms proposed in this paper are implemented using Answer Set Programming (ASP) (Gelfond and Lifschitz, 1988).

ASP is a declarative programming paradigm using logic programs under Answer Set semantics. In this paradigm the user provides a description of a problem and ASP works out how to solve the problem by returning answer sets corresponding to problem solutions. The existence of efficient solvers to generate the answers to the problems provided has increased the use of ASP in different domains of autonomous agents and multi-agent systems such as planning (Lifschitz, 2002) and normative reasoning (Cliffe et al., 2006; Panagiotidi et al., 2012b). Several action and planning languages such as event calculus (Kowalski and Sergot, 1986), $\mathcal{A}$ (and its descendants $\mathcal{B}$ and $\mathcal{C}$ (Gelfond and Lifschitz, 1998), Temporal Action Logics (TAL) (Doherty et al., 1998), have been implemented in ASP (Lee and Palla, 2012, 2014), indicating that ASP is appropriate for reasoning about actions. This provides motive and justification for an implementation of STRIPS (Fikes and Nilsson, 1971) that serves as the foundation of our model in ASP.

This paper is organised as follows. First we present a scenario in Section 2 which we use to illustrate the applicability of our approach. This is followed by the formal model and its semantics in Section 3. The computational

implementation of the model is provided in Section 4. After the discussion of related work in Section 5, we conclude in Section 6.

2. Illustrative Scenario

To illustrate our approach and motivate the normative practical reasoning model in the next section, we consider a scenario in which a software agent acts as a supervisory system in a disaster recovery mission and supports human decision-making in response to an emergency. The software agent's responsibility is to provide humans with different courses of actions available and to help humans decide on which course of actions to follow. In our scenario, the agent is to plan for a group of human actors who are in charge of responding to an emergency caused by an earthquake. The agent monitors the current situation (e.g., contamination of water, detection of shocks, etc.) and devises potential plans to satisfy goals set by human actors. In our scenario we assume the following two goals:

1. Running a hospital to help wounded people: this goal is fulfilled when medics are present to offer help and they have access to water and medicines.
2. Organising a survivors' camp: this is fulfilled when the camp's area is secured and a shelter is built.

We also assume the two following norms that the agent has to consider while devising plans to satisfy the goals above:

1. It is forbidden to built a shelter within 3 time units of detecting shocks. The cost of violating this norm is 5 units.
2. It is obligatory to stop water distribution for 2 time units once poison is detected in the water. The cost of violating this norm is 10 units.

The formulation of this scenario is provided in Appendix A.

3. A Model for Normative Practical Reasoning

In this research, we take STRIPS (Fikes and Nilsson, 1971) as the basis of our normative practical reasoning model. In STRIPS, a planning problem is defined in terms of an initial state, a goal state and a set of operators (e.g. actions). Each operator has a set of preconditions representing the circumstances/context in which the operator can be executed, and a set of postconditions that result from applying the operator. Any sequence of actions that satisfies the goal is a solution to the planning problem. In order

to capture the features of the normative practical reasoning problem that we are going to model, we extend the classical planning problem by:

(i) replacing atomic actions with *durative* actions: often the *nature* of the actions is non-atomic, which means that although executed atomically in a state, the system state in which they finish executing is not necessarily the same in which they started (Nunes et al., 1997). Refinement of atomic actions to durative actions reflects the real time that a machine takes to execute certain actions, which is also known as "real-time duration" of actions (Börger and Stärk, 2003).

(ii) Allowing a set of potentially inconsistent goals instead of the conventional single goal: the issue of planning for multiple goals distributed among distinct agents is addressed in collaborative planning. We address this issue for a single agent when handling multiple conflicting goals.

(iii) Factoring in norms: having made a case for the importance of norms in Section 1, we combine normative and practical reasoning. Just like goals, a set of norms is not necessarily consistent, making it potentially impossible for an agent to comply with all norms imposed on it.

A solution for a normative practical reasoning problem that features (i), (ii) and (iii) is any sequence of actions that satisfies at least one goal. The agent has the choice of violating or complying with norms triggered by execution of a sequence of actions, while satisfying its goals. However, there may be consequences either way that the agent has to foresee.

We explain the syntax and semantics of the model in Sections 3.2–3.6. First, however, we present the architecture of our envisaged system in the next section.

3.1. Architecture

The architecture, depicted in Figure 1, shows how re-planning can be considered if a plan in progress is interrupted due to a change in circumstances. This change can be the result of a change in the environment or unexpected actions of other agents in the system. As is customary in multi-agent systems, the agent will regularly check the viability of its plan. The frequency depends on the type of system the agent is operating in, the agent's commitment to its intentions, and the impact of re-computation on the agent's overall performance.

When an agent decides that re-planning is in order, it will take the state in which the plan is interrupted as the initial state for the new plan and its

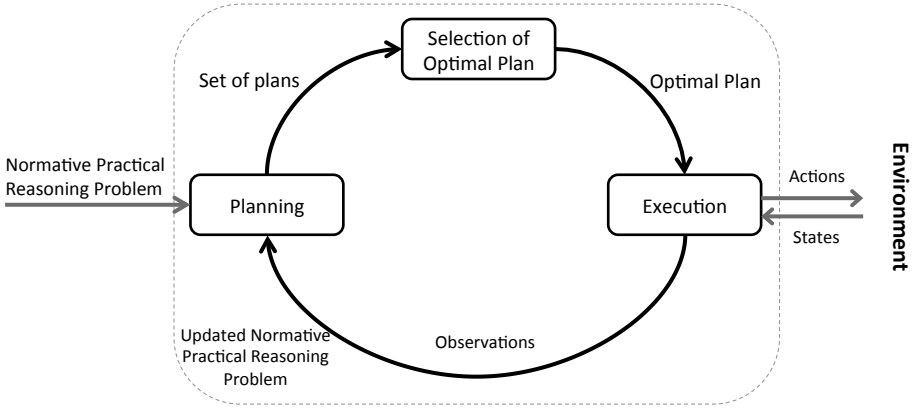

Figure 1: Overview of the System Architecture

current goal set as the goals to plan towards. The current goal set does not have to be the same as the goal set the original plan was devised for. Goals can already be achieved in the interrupted plan, previous goals may no longer be relevant and others may have been added. Even if the goals remain the same, the resulting new optimal plan might change due to changes in the state of the system. Similarly, there might be new norms imposed on the agent that will have to be considered in replanning.

We cater for agents which need to create their own individual plans. However, in doing so, in multi-agent scenarios agents will inevitably interact and interfere with each other's plans. The overview of Figure 1 will cater for this in two ways: (i) agents will notice the states being changed by other agents (a way of indirect communication) and (ii) the "Observations" will also contain interactions among agents (direct communication). Although we do not explore these ideas in the present paper, we envisage that the outcome of the indirect and direct interactions among agents could be represented as norms, establishing, for instance, that a particular agent, in the current context (*i.e.*, power or organisational relationships, description of capabilities and global goals, etc.), is forbidden or obliged to do something.

3.2. Syntax

We start this section by describing an extended STRIPS planning problem, defined in (Shams, 2016), that accommodates (i) durative actions; (ii) multiple goals and (iii) multiple norms.

Definition 1 (Normative Practical Reasoning Problem)**.** *A normative practical reasoning problem is a tuple* $P = (FL, \Delta, A, G, N)$ *where*

(i) FL *is a set of fluents;*
(ii) Δ *is the initial state;*
(iii) A *is a finite, non-empty set of durative actions for the agent;*
(iv) G *is the agent's set of goals;*
(v) N *is a set of norms.*

We describe in the ensuing sections each of the components of a normative practical reasoning problem.

3.2.1. Fluents and Initial State

FL is a set of domain fluents describing the domain the agent operates in. A literal l is a fluent or its negation i.e. $l = fl$ or $l = \neg fl$ for some $fl \in FL$. For a set of literals L, we define $L^+ = \{fl \,|\, fl \in L\}$ and $L^- = \{fl \,|\, \neg fl \in L\}$ to denote the set of positive and negative fluents in L respectively. L is well-defined if there exists no fluent $fl \in FL$ such that $fl, \neg fl \in L$, i.e., if $L^+ \cap L^- = \emptyset$.

The semantics of the normative practical reasoning problem is defined over a set of states Σ. A state $s \subseteq FL$ is determined by set of fluents that hold *true* at a given time, while the other fluents (those that are not present) are considered to be false. A state $s \in \Sigma$ satisfies fluent $fl \in FL$, denoted $s \models fl$, if $fl \in s$. It satisfies its negation $s \models \neg fl$ if $fl \notin s$. This notation can be extended to a set of literals as follows: set X is satisfied in state s, $s \models X$, when $\forall x \in X, s \models x$.

The set of fluents that hold at the initial state is denoted by $\Delta \subseteq FL$.

3.3. Durative Actions

The component A of our normative practical reasoning problem $P = (FL, \Delta, A, G, N)$ is a set of durative actions. A durative action has pre- and post-conditions. The effects of an action (as captured by its post-conditions) are not immediate, that is, it takes a non-zero period of time for the effects of an action to take place.

Definition 2 (Durative Actions)**.** *A durative action* a *is a tuple* $\langle pr, ps, d \rangle$ *where* pr *and* ps *are possibly empty and finite sets of well-defined literals representing respectively the pre- and post-conditions of the action, and* $d \in \mathbb{N}, d > 0$*, is the duration of the action.*

Given an action $a = \langle pr, ps, d \rangle$ we may also refer to its three components $pr(a), ps(a)$ and d_a. Moreover, we use $pr(a)^+$ and $pr(a)^-$ to refer to, respectively, the positive and negative literals in $pr(a)$; similarly, we have $ps(a)^+$ and $ps(a)^-$ to refer to respectively the positive and negative literals in $ps(a)$.

An action a can be executed in a state s if its preconditions hold in that state (i.e. $s \models pr(a)$). When modelling durative actions, there might be several states between the start and end state of the action, during which the action is said to be *in progress*. Some approaches take the view that it is sufficient for the preconditions of the action to hold at the start state and it does not matter whether they hold while the action is in progress (Knoblock, 1994), whereas others hold the view that the preconditions of an action should be satisfied while the action is in progress (Blum and Furst, 1997). Moreover, some planning languages, such as Planning Domain Description Language (PDDL) (Garrido et al., 2002; Fox and Long, 2003), distinguish between preconditions and those conditions that have to hold while the action is in progress. The latter conditions are referred to as *invariant conditions*. Having invariant conditions different from preconditions undoubtedly brings more expressiveness to the planning language; however it comes at the price of higher implementation complexity. In this paper we take the position that the invariant conditions are the same as preconditions, which implies that the preconditions have to be preserved throughout the execution of the action.

The postconditions of a durative action cause changes to the state s in which the action ends. These changes are: adding the positive postconditions $ps(a)^+$ to s and deleting the negative postconditions $ps(a)^-$ from s. Thus, for a state s in which action a ends, we have: $s \models ps(a)^+$ and $s \not\models ps(a)^-$.

Example 1. *Action "buildShelter" is available to the agents in the scenario described in Section 2. To build a shelter the agent has to secure and evacuate the area and there is no Shock detected – these are represented as preconditions areaSecured, evacuated and ¬ShockDetected, respectively. Once the shelter is built the area does not have to remain evacuated. This is represented as the positive postcondition shelterBuilt, while the negative postcondition of this action is ¬evacuated. In our scenario, we model this action as taking 4 units of time (that is, if it is executed in state s_k, it will end in state s_{k+4},).*

$$buildShelter = \left\langle \left\{ \begin{array}{c} areaSecured, \\ evacuated, \\ \neg ShockDetected \end{array} \right\}, \left\{ \begin{array}{c} shelterBuilt, \\ \neg evacuated \end{array} \right\}, 4 \right\rangle$$

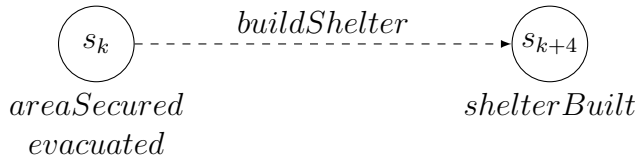

3.4. Goals

Goals identify the state of affairs in the world that an agent wants to achieve. Different types of goals and their characteristics have been classified in the literature (van Riemsdijk et al., 2008). *Achievement* goals are the most common type of goals modelled in the agent literature and have therefore received the most attention (van Riemsdijk et al., 2008; de Boer et al., 2002; Nigam and Leite, 2006; van Riemsdijk et al., 2002). goals for the purpose of this research are achievement goals.

We define below the elements of the set G of $P = (FL, \Delta, A, G, N)$.

Definition 3 (Goals)**.** *A goal $g \in G$ is the pair $\langle r, v \rangle$, where r is a possibly empty and finite set of well-defined literals representing the goal requirements, and $v \in \mathbb{N}, v > 0$, represents the utility/gain for achieving the goal.*

Goal g's requirements and value are denoted as $r(g)$ and $v(g)$, respectively.

Example 2. *The goals from our illustrative scenario are formulated as below:*

$$runningHospital = \left\langle \left\{ \begin{array}{c} medicsPresent, \\ waterSupplied, \\ medicinesSupplied \end{array} \right\}, 25 \right\rangle$$

$$organiseSurvivorCamp = \left\langle \left\{ \begin{array}{c} areaSecured, \\ shelterBuilt \end{array} \right\}, 18 \right\rangle$$

3.5. Norms

In this section we specify what we refer to as a norm in this work. In order to provide a context for the norm specification we explain how our norm specification corresponds to the five elements identified by Criado (2012) that distinguish norm specification languages.

1. Deontic Operators: We model a permissive society in which the agent has complete knowledge of the domain of actions available. Everything is permitted unless it is explicitly prohibited. The role of obligation is to motivate the agent to execute a specific action and the role of prohibition is to inhibit the agent from executing a particular action.

2. Controls: Controls determine whether the deontic propositions operate on actions, states or both. In this work we focus on action-based norms.
3. Enforcement Mechanisms: We use the enforcement mechanism proposed by Shams et al. (2015) that is a combination of utility-based (e.g., Oren et al. (2011); Panagiotidi et al. (2012a)) and pressure-based (López y López et al., 2005) compliance methods.
4. Conditional Expressions: Similar to the control element, we use actions as conditional expressions. In other words, the norm condition is an action that once executed, the agent is obliged to or prohibited from executing the action that the norm targets.
5. Temporal Constraints: temporal constraints can be used to express norm activation, termination, deadline, etc. The temporal constraint we specify here is concerned with the deadline. The agent is expected to comply with an obligation (execute a certain action) or a prohibition (refrain from executing a specific action) before some deadline.

Having explained the properties of our norm specification language, we now define the element N of problem $P = (FL, \Delta, A, G, N)$. N denotes a set of conditional norms to which the agent is subject:

Definition 4 (Norms)**.** *N is a set of norms, each of which is a tuple of the form $n = \langle d_o, a_{con}, a_{sub}, dl, c\rangle$, where*

- *$d_o \in \{o, f\}$[2] is the deontic operator determining the type of norm, which can be an obligation or a prohibition;*
- *$a_{con} \in A$ is the durative action (cf. Def. 2) that activates the norm;*
- *$a_{sub} \in A$ is the durative action (cf. Def. 2) that is the target of the obligation or prohibition;*
- *$dl \in \mathbb{N}$ is the norm deadline relative to the activation condition, which is the completion of the execution of the action a_{con}; and*
- *$c \in \mathbb{N}$ is the penalty cost that will be applied if the norm is violated. $c(n)$ denotes the penalty cost of norm n.*

An obligation norm states that executing action a_{con} obliges the agent to start/start and end the execution of a_{sub} within dl time units of the end of execution of a_{con}. Such an obligation is complied with if the agent starts or

[2]The symbols o and f are normally represented as respectively **O** and **F** in the Deontic logic literature. However we have used lower case letters to make these consistent with our implementation in the next section. Capital letters in the implementation language are reserved for variables.

starts and ends executing a_{sub} before the deadline and is violated otherwise. A prohibition norm expresses that executing action a_{con} prohibits the agent from starting or starting and ending the execution of a_{sub} within dl time units of the end of execution of a_{con}. Such a prohibition is complied with if the agent does not start or does not start and end executing a_{sub} before the deadline and is violated otherwise.

Example 3. *The norms in the illustrative scenario are formulated as below:*

$$n_1 = \langle f, detectShock, buildShelter, 3, 5\rangle$$
$$n_2 = \langle o, detectPoison, stopWater, 2, 10\rangle$$

A norm can be activated multiple times in a sequence of action, generating different instances of the original norm. To make sure different instances are dealt with uniquely, we define *instantiated norms*. In each instance the deadline is updated relative to the end of execution of the action that is the condition of the norm.

Definition 5 (Instantiated Norm)**.** *An instantiation of norm* $n = \langle d_o, a_{con}, a_{sub}, dl, c\rangle$ *is denoted as* $n_{ins} = \langle d_o, a_{con}, a_{sub}, dl_{ins}, c\rangle$ *where* $dl_{ins} = dl + t_{a_{con}} + d_{a_{con}}$. $t_{a_{con}}$ *is when action* a_{con} *is executed and* $d_{a_{con}}$ *is the duration of* a_{con}.

We also denote an instantiation of a norm n_i as n'_i.

Example 4. *Assume that in some sequence of action detectShock is executed at time* 3 *(i.e.* $t_{a_{con}} = 3$*) and that the duration of this action is* 1 *(i.e.* $d_{a_{con}} = 1$*). The instantiated version of norm*

$$n_1 = \langle f, detectShock, buildShelter, 3, 5\rangle$$

in this sequence of actions is

$$n'_1 = \langle f, detectShock, buildShelter, 7, 5\rangle$$

Where dl_{ins} *is calculated based on Def. 5*

3.6. Semantics

Having explained the syntax of the model, we now focus on the semantics. To this end, we need to describe given a normative practical reasoning problem $P = (FL, \Delta, A, G, N)$:

(i) What the possible courses of action for the agent are and what properties each course of action has. Properties are defined in terms of goals that a sequence of action satisfies, norms it complies with and the norms it violates. This item is discussed in Section 3.6.1.

(ii) What different type of conflicts the agent can experience while trying to satisfy its goals and comply with the norms to which it is subject. In Section 3.6.2 we explore this item.

(iii) What identifies a sequence of actions as a solution/plan for problem P. Plans are defined in Section 3.6.3.

3.6.1. Sequences of Actions and their Properties

Let $P = (FL, \Delta, A, G, N)$ be a normative practical reasoning problem. Also let $\pi = \langle (a_0, 0), \cdots, (a_n, t_{a_n}) \rangle$ with $a_i \in A$ and $t_{a_i} \in \mathbb{Z}^+$ be a sequence of actions a_i executed at time t_{a_i}. The pair (a_i, t_{a_i}) reads as action a_i is executed at time $t_{a_i} \in \mathbb{Z}^+$ s.t. $\forall i < j, t_{a_i} < t_{a_j}$. The total duration of a sequence of actions, $Makespan(\pi)$, is defined in Equation 1.

$$Makespan(\pi) = max(t_{a_i} + d_{a_i}) \tag{1}$$

Actions in a sequence can be executed concurrently but they cannot start at the same time. This is because the planning problem is defined for a single agent and a single agent is not typically assumed to be able to start two actions at the exact same instant. Also actions in the sequence should not have concurrency conflicts, which are defined below. In our presentation we need to check for the occurrence of specific pairs in a sequence of actions π, and we thus define the operator "$\widehat{\in}$" as

$$(a, t_a) \,\widehat{\in}\, \pi \text{ iff } \begin{cases} \pi = \langle (a, t_a), \ldots, (a_n, t_n) \rangle \text{ or} \\ \pi = \langle (a_0, 0), \ldots, (a, t_a), \ldots, (a_n, t_n) \rangle \text{ or} \\ \pi = \langle (a_0, 0), \ldots, (a, t_a) \rangle \end{cases}$$

Temporary conflicts prevent the agent from executing two actions under specific constraints, the most common one of which is time. Conflicts caused by time, known as *concurrency conflicts* between actions, prevent actions from being executed in an overlapping period of time. Blum and Furst (1997) define that two actions a_i and a_j cannot be executed concurrently, if at least one of the following holds:

1. The preconditions of a_i and a_j contradict each other:

$$\exists r \in pr(a_i) \text{ s.t. } \neg r \in pr(a_j) \quad or$$

$$\exists \neg r \in pr(a_i) \text{ s.t. } r \in pr(a_j)$$

2. The postconditions of a_i and a_j contradict each other:

$$\exists r \in ps(a_i)^+ \text{ s.t. } \neg r \in ps(a_j)^- \quad or$$

$$\exists \neg r \in ps(a_i)^- \text{ s.t. } r \in ps(a_j)^+$$

3. The postconditions of a_i contradict the preconditions of a_j:

$$\exists r \in ps(a_i)^+ \text{ s.t. } \neg r \in pr(a_j) \quad or$$

$$\exists \neg r \in ps(a_i)^- \text{ s.t. } r \in pr(a_j)$$

4. The preconditions of a_i are contradicted by the postconditions of a_j:

$$\exists r \in pr(a_i) \text{ s.t. } \neg r \in ps(a_j)^- \quad or$$

$$\exists \neg r \in pr(a_i) \text{ s.t. } r \in ps(a_j)^+$$

We summarise the four conditions above in Definition 6, where we define what are referred to as conflicting actions in the remainder of this work.

Definition 6 (Conflicting Actions)**.** *Actions a_i and a_j have a concurrency conflict if the pre- or post-conditions of a_i contradict the pre- or post-conditions of a_j. The set of conflicting actions is denoted as cf_{action}:*

$$cf_{action} = \left\{ (a_i, a_j) \;\middle|\; \begin{array}{l} \exists r \in pr(a_i) \cup ps(a_i)^+, \neg r \in pr(a_j) \cup ps(a_j)^- \\ \quad or \\ \exists \neg r \in pr(a_i) \cup ps(a_i)^-, r \in pr(a_j) \cup ps(a_j)^+ \end{array} \right\} \quad (2)$$

Example 5. *Assume action evacuate with the following specification:*

$$evacuate = \left\langle \left\{ \begin{array}{c} populated, \\ ShockDetected, \end{array} \right\}, \left\{ \begin{array}{c} evacuated, \\ \neg populated \end{array} \right\}, 5 \right\rangle$$

The pre- and post-conditions of this action are inconsistent with the pre- and post-conditions of action buildShelter defined in Example 1:

$$buildShelter = \left\langle \left\{ \begin{array}{c} areaSecured, \\ evacuated, \\ \neg ShockDetected \end{array} \right\}, \left\{ \begin{array}{c} shelterBuilt, \\ \neg evacuated \end{array} \right\}, 4 \right\rangle$$

Therefore, these two actions cannot be executed concurrently. However, action evacuate effectively contributes to the preconditions of buildShelter, which means they can indeed be executed consecutively.

Definition 7 (Sequence of States)**.** *Let* $\pi = \langle (a_0, 0), \cdots, (a_n, t_{a_n}) \rangle$ *be a sequence of actions such that* $\nexists (a_i, t_{a_i}), (a_j, t_{a_j}) \in \pi$ *s.t.* $t_{a_i} \leq t_{a_j} < t_{a_i} + d_{a_i}, (a_i, a_j) \in cf_{action}$ *and let* $m = Makespan(\pi)$*. The execution of a sequence of actions* π *from a given starting state* $s_0 = \Delta$ *brings about a sequence of states* $S(\pi) = \langle s_0, \cdots s_m \rangle$ *for every discrete time interval from* 0 *to* m*.*

The transition relation between states is given by Def. 8. If action a_i ends at time k, state s_k results from removing delete post-conditions and adding add post-conditions of action a_i to state s_{k-1}. If there is no action ending at s_k, the state s_k remains the same as s_{k-1}. We first define A_k as the set of action/time pairs such that the actions end at some specific state s_k:

$$A_k = \{(a_i, t_{a_i}) \in \pi \,|\, k = t_{a_i} + d_{a_i}\} \tag{3}$$

Note that s_k is always well-defined since two actions with inconsistent post-conditions, according to Def. 6 belong to cf_{action} so they cannot be executed concurrently and thus they cannot end at the same state.

Definition 8 (State Transition)**.** *Let* $\pi = \langle (a_0, 0), \cdots, (a_n, t_{a_n}) \rangle$ *and let* $S(\pi) = \langle s_0, \cdots s_m \rangle$ *be the sequence of states brought about by* π*:*

$$\forall k > 0 : \; s_k = \begin{cases} (s_{k-1} \setminus (\bigcup_{a \in A_k} ps(a)^-) \cup \bigcup_{a \in A_k} ps(a_i)^+ & A_k \neq \emptyset \\ s_{k-1} & A_k = \emptyset \end{cases} \tag{4}$$

We now turn our attention to the properties of each sequence of actions.

Definition 9 (Goal Satisfaction)**.** *Goal requirements should hold in order to satisfy the goal. A sequence of actions* $\pi = \langle (a_0, 0), \cdots (a_n, t_{a_n}) \rangle$ *satisfies goal* g *if there is at least one state* $s_k \in S(\pi)$ *that satisfies the goal:*

$$\pi \models r(g) \textit{ iff } \exists \; s_k \in S(\pi) \textit{ s.t. } s_k \models r(g) \tag{5}$$

The set of goals satisfied by π is denoted as G_π:

$$G_\pi = \{g \,|\, \pi \models r(g)\} \tag{6}$$

Definition 10 (Activated Norms)**.** *A norm* $n = \langle d_o, a_{con}, a_{sub}, dl, c \rangle$ *is instantiated in a sequence of actions* $\pi = \langle (a_0, 0), \cdots, (a_n, t_{a_n}) \rangle$ *if its activation condition* a_{con} *belongs to the sequence of actions. Let* N_π *be the set of instantiations of various norms in* π *defined in Equation 7. Note that* dl_{ins} *is calculated based on Definition 5.*

$$N_\pi = \{\langle d_o, a_{con}, a_{sub}, dl_{ins}, c \rangle \mid \langle d_o, a_{con}, a_{sub}, dl, c \rangle \in N, (a_{con}, t_{a_{con}}) \,\widehat{\in}\, \pi\} \tag{7}$$

Definition 11 (Obligation Compliance)**.** *A sequence of actions* $\pi = \langle (a_0, 0), \cdots, (a_n, t_{a_n}) \rangle$ *complies with an obligation* $n = \langle o, a_{con}, a_{sub}, dl_{ins}, c \rangle$ *if* a_{con} *is executed in* π *and* a_{sub}*, starts (cf. Eq. 8) or starts and ends (cf. Eq. 9) within the period when the condition holds and when the deadline expires:*

$$\pi \models n \text{ iff } (a_{con}, t_{a_{con}}), (a_{sub}, t_{a_{sub}}) \mathrel{\widehat{\in}} \pi \text{ s.t. } t_{a_{sub}} \in [t_{a_{con}} + d_{a_{con}}, dl_{ins}) \quad (8)$$

$$\pi \models n \text{ iff } (a_{con}, t_{a_{con}}), (a_{sub}, t_{a_{sub}}) \mathrel{\widehat{\in}} \pi \text{ s.t. } [t_{a_{sub}}, t_{a_{sub}} + d_{a_{sub}}] \subseteq [t_{a_{con}} + d_{a_{con}}, dl_{ins}) \quad (9)$$

Definition 12 (Obligation Violation)**.** *A sequence of actions* $\pi = \langle (a_0, 0), \cdots, (a_n, t_{a_n}) \rangle$ *violates obligation* $n_{ins} = \langle o, a_{con}, a_{sub}, dl_{ins}, c \rangle$ *if* a_{con} *is executed in* π*, but* a_{sub} *does not start (Equation 10), or does not start and end (Equation 11) in the period between the state when the condition holds and when the deadline expires.*

$$\pi \not\models n \text{ iff } (a_{con}, t_{a_{con}}) \mathrel{\widehat{\in}} \pi, (a_{sub}, t_{a_{sub}}) \mathrel{\widehat{\notin}} \pi \text{ s.t. } t_{a_{sub}} \in [t_{a_{con}} + d_{a_{con}}, dl_{ins}) \quad (10)$$

$$\pi \not\models n \text{ iff } (a_{con}, t_{a_{con}}) \mathrel{\widehat{\in}} \pi, (a_{sub}, t_{a_{sub}}) \mathrel{\widehat{\notin}} \pi \text{ s.t. } [t_{a_{sub}}, t_{a_{sub}} + d_{a_{sub}}] \subseteq [t_{a_{con}} + d_{a_{con}}, dl_{ins}) \quad (11)$$

Example 6. *Let there be the following instantiated version*

$$n_2' = \langle o, detectPoison, stopWater, 8, 10 \rangle$$

of norm

$$n_2 = \langle o, detectPoison, stopWater, 2, 10 \rangle$$

from Example 3. The compliance period for this norm in displayed in the figure below. According to Def. 11 in its Eq. 8, if $t_{stopWater}$ *belongs to this period, this norm instance is complied with; otherwise, according to Def. 12 in its Eq. 10, the norm is violated. This is illustrated by the following diagram:*

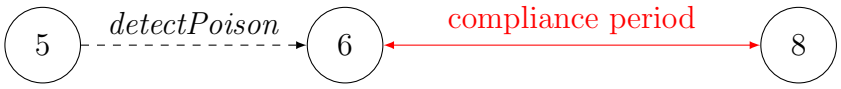

Example 7. *Let there be the following instantiated version*

$$n'_3 = \langle o, detectEarthquake, blockMainRoad, 7, 12\rangle$$

of norm

$$n_3 = \langle o, detectEarthquake, blockMainRoad, 5, 12\rangle$$

which obliges the agent to have blocked the main road within 5 units of time after detecting an earthquake. Since the post-conditions of action blockMainRoad are brought about at the end of its execution, according to Def. 11 (Eq. 9), this norm is complied with if blockMainRoad starts and ends between time points 2 and 7. Otherwise, according to Def. 12 (Eq. 11) this norm is violated.

Definition 13 (Prohibition Compliance)**.** *A sequence of actions* $\pi = \langle (a_0, 0), \cdots, (a_n, t_{a_n})\rangle$ *complies with prohibition* $n = \langle f, a_{con}, a_{sub}, dl_{ins}, c\rangle$ *if* a_{con}*, is executed and* a_{sub}*, does not start (Eq. 12) or does not start and end (Eq. 13) in the period when the condition holds and the deadline expires. Formally:*

$$\pi \models n \text{ iff } (a_{con}, t_{a_{con}}) \,\widehat{\in}\, \pi, (a_{sub}, t_{a_{sub}}) \,\widehat{\notin}\, \pi \text{ s.t.}$$
$$t_{a_{sub}} \in [t_{a_{con}} + d_{a_{con}}, dl_{ins}) \quad (12)$$

$$\pi \models n \text{ iff } (a_{con}, t_{a_{con}}) \,\widehat{\in}\, \pi, (a_{sub}, t_{a_{sub}}) \,\widehat{\notin}\, \pi \text{ s.t.}$$
$$[t_{a_{sub}}, t_{a_{sub}} + d_{a_{sub}}] \subseteq [t_{a_{con}} + d_{a_{con}}, dl_{ins}) \quad (13)$$

Definition 14 (Prohibition Violation)**.** *A sequence of actions* $\pi = \langle (a_0, 0), \cdots, (a_n, t_{a_n})\rangle$ *violates prohibition* $n = \langle f, a_{con}, a_{sub}, dl_{ins}, c\rangle$ *iff* a_{con}*, has occurred and* a_{sub} *starts (Eq. 14) or starts and ends (Eq. 15) in the period between when the condition holds and when the deadline expires. Formally:*

$$\pi \not\models n \text{ iff } (a_{con}, t_{a_{con}}), (a_{sub}, t_{a_{sub}}) \,\widehat{\in}\, \pi \text{ s.t. } t_{a_{sub}} \in [t_{a_{con}} + d_{a_{con}}, dl_{ins}) \quad (14)$$

$$\pi \not\models n \text{ iff } (a_{con}, t_{a_{con}}), (a_{sub}, t_{a_{sub}}) \,\widehat{\in}\, \pi \text{ s.t.}$$
$$[t_{a_{sub}}, t_{a_{sub}} + d_{a_{sub}}] \subseteq [t_{a_{con}} + d_{a_{con}}, dl_{ins}) \quad (15)$$

Example 8. *Let there be the following instantiated version*

$$n_1' = \langle f, detectShock, buildShelter, 6, 5 \rangle$$

of norm

$$n_1 = \langle f, detectShock, buildShelter, 3, 5 \rangle$$

presented in Ex. 3. The compliance period for this norm in displayed in the figure below. According to Def. 14 (Eq. 14), if $t_{buildShelter}$ belongs to this period, this norm instance is violated; otherwise, according to Def. 13 (Eq. 12), it is complied with.

Example 9. *Let there be the following instantiated version*

$$n_4' = \langle f, detectEarthquake, blockMainRoad, 7, 12 \rangle$$

of norm

$$n_4 = \langle f, detectEarthquake, blockMainRoad, 5, 12 \rangle$$

which forbids the agent from blocking the main road within 5 units of time after detecting an earthquake. Since the post-conditions of action blockMainRoad are brought about at the end of its execution, according to Def. 13 (Eq. 13), this norm is violated if blockMainRoad starts and ends between time points 2 and 7. Otherwise, according to Def. 14 (Eq. 15) this norm is complied with. This is illustrated by the diagram below.

The set of norms complied with and violated in π are denoted as $N_{cmp(\pi)}$ and $N_{vol(\pi)}$ respectively, and defined as follows:

$$N_{cmp(\pi)} = \{n_{ins} \in N_\pi \mid \pi \models n_{ins}\} \tag{16}$$

$$N_{vol(\pi)} = \{n_{ins} \in N_\pi \mid \pi \not\models n_{ins}\} \tag{17}$$

To make sure there are no norms pending at $m = Makespan(\pi)$, we assume that the norm deadlines are smaller than m. Therefore, all the activated norms in π are either complied with or violated by time m:

$$N_\pi = N_{cmp(\pi)} \cup N_{vol(\pi)} \tag{18}$$

Alternatively, it could be assumed that the norms that are pending (i.e. neither violated nor complied with) at m are all considered as complied with or violated.

3.6.2. Conflict

In the previous section we defined when a sequence of actions satisfies a goal, complies with or violates a norm. A possible cause for not satisfying a certain goal is the conflict between the goal and another goal or norm. Likewise, violating a norm could help in reaching a goal or complying with another norm. In this work we do not concern ourselves directly with detecting or resolving conflicts, instead, we focus on the consequences of such conflicts on the agent behaviour. To make this work self-contained, however, we briefly review the causes of conflicts between goals, between norms and between goals and norms. We leave for future work the provision of agents with the capability to reason about the plans and consequently *inferring* the conflict between goals, between norms and between goals and norms.

An agent may pursue multiple goals or desires at the same time and it is likely that some of these goals conflict (van Riemsdijk et al., 2002; Nigam and Leite, 2006; Pokahr et al., 2005; Thangarajah et al., 2003; van Riemsdijk et al., 2009). Conflict between the agent's goals or desires, especially for BDI agents, has been addressed by several authors. Hulstijn and van der Torre (2004) describe two goals as conflicting if achieving them requires taking two conflicting actions, where conflicting actions are encoded using integrity constraints. Rahwan and Amgoud (2006) on the other hand, define two desires as conflicting if the sets of beliefs that supports the achievement of desires are contradictory. Like Rahwan and Amgoud (2006), Broersen et al. (2002) argue that for a set of goals not to be conflicting, a consistent mental attitude (e.g. beliefs and norms) is required. Some (*e.g.*, (Toniolo, 2013)) have adopted a static view on goal conflict, in which conflicting goals are mutually-exclusive, hence impossible to satisfy in the same plan regardless of the order or choice of actions in the plan. Limited and bounded resources (*e.g.* time, budget, etc.) are debated as another cause of conflict between goals (Thangarajah et al., 2002).

López y López et al. (2005) discuss conflict between goals and norms in terms of goals being hindered by norms or vice-versa. The same applies to the approach offered by Modgil and Luck (2008), suggesting a mechanism to resolve the conflicts between desires and normative goals. In this approach, norms are represented as system goals that may conflict with an agent's goals or desires. Social goals and individual goals do not need to conflict directly. Instead, conflict arises from the reward or punishment of complying with or violating a norm that may facilitate or hinder some of the agent's individual

goals. Shams et al. (2015) identify the conflict between norms and goals as follows. When an obligation forces the agent to take an action that has postconditions that are inconsistent with the requirements of a goal, they may come into conflict. On the other hand, when an action is prohibited, and the postconditions of that action contribute to a goal, they may conflict.

Conflict between norms have been studied in multi-agent systems (*e.g.*, Vasconcelos et al. (2009)) as well as other domains such as legal reasoning (*e.g.*, (Sartor, 1992)). When faced with conflicting norms, the agent cannot comply with both of them and hence one of the norms is violated. In terms of action-based norms, Shams et al. (2015) define two obligations conflicting if they oblige the agent to take two conflicting actions (cf. Def. 6) in an overlapping time interval. Likewise, an obligation and a prohibition cause conflict if they oblige and forbid the agent to execute the same action in an overlapping time interval.

Having defined sequences of actions and the properties and conflicts they can involve, we can now define which sequences of action can be identified as plans in the next section.

3.6.3. Plans

In classical planning a sequence of actions $\pi = \langle (a_0, 0), \cdots, (a_n, t_n) \rangle$ is identified as a plan if all the fluents in the initial state, do hold at time 0 and for each i, the preconditions of action a_i hold at time t_{a_i}, and the goal of planning problem is satisfied in time m, where $m = \mathit{Makespan}(\pi)$. However, extending the conventional planning problem by multiple potentially conflicting goals and norms requires defining extra conditions in order to make a sequence of actions a plan and a solution for P. In what follows, we define what is required to identify a sequence of actions as a plan.

Definition 15 (Plan)**.** *A sequence of actions* $\pi = \langle (a_0, 0), \ldots, (a_n, t_{a_n}) \rangle$ *s.t.* $\nexists (a_i, t_{a_i}), (a_j, t_{a_j}) \widehat{\in} \pi$ *s.t.* $t_{a_i} \leq t_{a_j} < t_{a_i} + d_{a_i}, (a_i, a_j) \in cf_{action}$ *is a plan for the normative practical reasoning problem* $P = (FL, \Delta, A, G, N)$ *if the following conditions hold:*

- *fluents in* Δ *(and only those fluents) hold in the initial state:* $s_0 = \Delta$
- *the preconditions of action* a_i *holds at time* t_{a_i} *and throughout the execution of* a_i*:* $\forall k \in [t_{a_i}, t_{a_i} + d_{a_i}), s_k \models pr(a_i)$
- *plan* π *satisfies a non-empty subset of goals:* $G_\pi \neq \emptyset$

The utility of a plan π is defined by deducting the penalty costs of violated norms from the value gain of satisfying goals (Equation 19). The set of

optimal plans, *Opt*, are those plans that maximise the utility.

$$\mathit{Utility}(\pi) = \sum_{g_i \in G_\pi} v(g_i) - \sum_{n_j \in N_{vol(\pi)}} c(n_j) \tag{19}$$

Examples of calculating the utility of plans are in Appendix B.

The set *Opt* is empty only if there are no plans for the planning problem. Otherwise, the utility function is guaranteed to terminate and find the optimal plans and hence populate the set *Opt*.

4. Implementation

In this section, we demonstrate how a normative practical reasoning problem $P = (FL, \Delta, A, G, N)$ (cf. Def. 1), can be implemented. Our implementation should be seen as a proof of concept that provides a computational realisation of all aspects of the formal model. We use Answer Set Programming (ASP) (Gelfond and Lifschitz, 1988) to propose such an implementation. Recent work on planning in ASP (To et al., 2015) demonstrates that in terms of general planners ASP is a viable competitor. The Event Calculus (EC) (Kowalski and Sergot, 1986) forms the basis for the implementation of some normative reasoning frameworks, such as those of Alrawagfeh and Meneguzzi (2014) and Artikis et al. (2009). Our proposed formal model is independent of language and could be translated to EC and hence to a computational model. However, the one-step translation to an ASP is preferred because the formulation of the problem is much closer to a computational model, thus there is a much narrower conceptual gap to bridge. Furthermore, the EC implementation language is Prolog, which although syntactically similar to ASP, suffers from non-declarative features such as clause ordering affecting the outcome and the cut ("!") operator, jeopardising its completeness. Also, its query-based nature which focuses on one query at a time, makes it cumbersome to reason about all plans.

In what follows, we provide a brief introduction to ASP in Section 4.1, followed by the mapping of normative practical reasoning problem $P = (FL, \Delta, A, G, N)$ (cf. Def. 1) into ASP in Section 4.2. In the latter section we show how P is mapped into an answer set program such that there is a one to one correspondence between solutions for the problem and the answer sets of the program. The mapping itself is provided in Figure 2. The explanation of the mapping is presented in Sections 4.2.1–4.2.6, with cross references to the code fragments listed in Figure 2.

4.1. Answer Set Programming

ASP is a declarative programming paradigm using logic programs under Answer Set semantics (Lifschitz, 2008). Like all declarative paradigms it has the advantage of describing the constraints and the solutions rather than the writing of an algorithm to find solutions to a problem. A variety of programming languages for ASP exist, and we use *AnsProlog* (Baral, 2003). There are several efficient solvers for *AnsProlog*, of which CLINGO (Gebser et al., 2011) and DLV (Eiter et al., 1999) are currently the most widely used.

The basic components of *AnsProlog* are *atoms* that are constructs to which one can assign a truth value. An atom can be negated, adopting *negation as failure* (naf), which establishes that a negated atom not a is true if there is no evidence to prove a. *Literals* are atoms a or negated atoms not a (referred to as naf-literals). Atoms and literals are used to create rules of the general form "$a :- b_1, ..., b_m, \text{not } c_1, ..., \text{not } c_n$." where a, b_i and c_j are atoms. Intuitively, a rule means that *if all atoms b_i are known/true and no atom c_j is known/true, then a must be known/true.* We refer to a as the *head of the rule* and $b_1, ..., b_m, \text{not } c_1, ..., \text{not } c_n$ as the *body of the rule.* A rule with an empty body is called a *fact* and a rule with an empty head is called a *constraint*, indicating that no solution should be able to satisfy the body. Another type of rules are called choice rules and are denoted as $l\{h_0, \cdots, h_k\}u :- l_1, \cdots, l_m, not\ \ l_{m+1}, \cdots, not\ l_n.$, in which h_is and l_is are atoms. l and u are integers and the default values for them are 0 and 1, respectively. A choice rule is satisfied if the number of atoms belonging to $\{h_0, \cdots, h_k\}$ that are true/known is between the lower bound l and upper bound u. A *program* is a set of rules representing their conjunction. The semantics of *AnsProlog* is defined in terms of *answer sets*, i.e. assignments of true and false to all atoms in the program that satisfy the rules in a minimal and consistent fashion. A program may have zero or more answer sets, each corresponding to a solution. We refer to Appendix D and Baral (2003) for a formal treatment of the semantics of ASP.

4.2. Translating the Normative Practical Reasoning Problem into ASP

Prior to describing the mapping (Figure 2) of normative practical reasoning problem $P = (FL, \Delta, A, G, N)$ into ASP, we list the atoms that we use in the mapping in Table 1. The description of the mapping is presented in the following sections with references to Figure 2. Note that variables are in capitals and grounded variables are in small italics.

Table 1: Atoms and their Intended Meaning

Atom	**Intended meaning**
`state(S)`	`S` is a state
`holdsat(F,S)`	fluent `F` holds in state `S`
`terminated(F,S)`	fluent `F` terminates in state `S`
`action(A,D)`	`A` is an action with duration `D`
`executed(A,S)`	action `A` is executed in state `S`
`pre(A,S)`	the preconditions of action `A` hold in state `S`
`inprog(A,S)`	action `A` is in progress in state `S`
`goal(G,V)`	`G` is a goal with value `V`
`satisfied(G,S)`	goal `G` is satisfied in state `S`
`norm(N,C)`	`N` is a norm with the violation cost of `C`
`cmp(o|f(N,S1,A,DL),S2)`	the instantiated version of obligation or prohibition `N` in state `S1`, is complied with in state `S2`
`vol(o|f(N,S1,A,DL),S2)`	the instantiated version of obligation or prohibition `N` in state `S1`, is violated in state `S2`
`value(TV)`	`TV` is the sum of values of goals satisfied in a plan
`cost(TC)`	`TC` is the sum of violation costs of norms violated in a plan
`utility(U)`	`U` is the utility of a plan as the difference of values of goals satisfied and norms violated in the plan

4.2.1. States

In Section 3.6 we described the semantics $P = (FL, \Delta, A, G, N)$ over a set of states. The facts produced by line 1 provide the program with all available states for plans of maximum length q. Currently, the length of the plans needs to be determined experimentally. We plan to automate this using incremental features of ASP solver CLINGO4 (Gebser et al., 2011). Line 2 encodes that the initial fluents, ($x \in \Delta$) need to hold at state 0 which is achieved by the facts `holdsat(`x`,0)`. Fluents are inertial, they continue to hold unless they are terminated. Inertia is encoded in lines 3–4. Termination is modelled through the predicate `terminated(X,S)`.

4.2.2. Actions

This section describes the details of encoding of actions. Each durative action is encoded as `action(`a, d`)` (line 5), where a is the name of the action and d is its duration. The preconditions $pr(a)$ of action a hold in state s if $s \models pr(a)$. This is expressed in line 6 using atom `pre(`a`,S)` In order to make the coding more readable we introduce the shorthand `EX(X,S)`, where `X` is a set of fluents that should hold at state `S`. For all $x \in$ `X`, `EX(X,S)` is translated into `holdsat(`x`,S)` and for all $\neg x \in$ `X`, `EX(`$\neg$`X,S)` is translated into `not EX(`x`,S)` using negation as failure.

The agent has the choice to execute any of its actions in any state. This is expressed in the choice rule in line 7. Since no lower or upper bound is given for {`executed(A,S)`}, the default value of `0{executed(A,S)}1` is implied, meaning that the agent has the choice of whether or not to execute an action. Following the approach in Blum and Furst (1997), we assume that the preconditions of a durative action should be preserved when it is in progress. We first encode the description of an action in progress, followed by ruling out the possibility of an action being in progress in the absence of its preconditions. A durative action is in progress, `inprog(A,S)`, from the state in which it starts up to the state in which it ends (lines 8–9). Line 10, rules out the execution of an action, when the preconditions of the action do not hold during its execution. A further assumption made is that the agent cannot start two actions at exactly the same time (line 11–12). Once an action starts, the result of its execution is reflected in the state where the action ends. This is expressed through (i) lines 13–14 that allow the add postconditions of the action to hold when the action ends, and (ii) line 15–16 that allow the termination of the delete postconditions. Termination takes place in the state *before* the end state of the action: the reason for this is the inertia of fluents that was expressed in lines 3–4. Thus delete post-conditions of an action are terminated in the state before the end state of the action, so that they will not hold in the following state, in which the action ends (i.e. they are deleted from the state).

4.2.3. Goals and Norms

Line 17 encodes goal g with value v as a fact. Goal g is satisfied in state s if $s \models g$. This is expressed in line 18, where g^+ and g^- are the positive and negative literals in the set g.

For the norms we note that, following Definitions 11–14, compliance and violation of a norm can be established based on the start state of action's

Creating states: $\forall\, k \in [0, q]$

```
1  state(k).
```

Setting up the initial state: $\forall\, x \in \Delta$

```
2  holdsat(x,0).
```

Rule for fluent inertia

```
3  holdsat(X,S2) :- holdsat(X,S1), not terminated(X,S1),
4                   state(S1), state(S2), S2=S1+1.
```

Creating the actions and their preconditions: $\forall a \in A,\, a = \langle pr, ps, d\rangle$

```
5  action(a,d).
6  pre(a,S) :- EX(pr(a)^+,S), not EX(pr(a)^-,S), state(S).
```

Common constraints on action execution

```
7  {executed(A,S)} :- action(A,D), state(S).
8  inprog(A,S2) :- executed(A,S1), action(A,D),
9                  state(S1), state(S2), S1<=S2, S2<S1+D.
10 :- inprog(A,S), action(A,D), state(S), not pre(A,S).
11 :- executed(A1,S), executed(A2,S), A1!=A2,
12    action(A1,D1), action(A2,D2), state(S).
```

Adding positive postconditions of actions: $ps(a)^+ = X \Leftrightarrow \forall x \in X \cdot$

```
13 holdsat(x,S2) :- executed(a,S1), action(a,d),
14                  state(S1), state(S2), S2=S1+d.
```

Terminating negative post conditions of actions: $ps(a)^- = X \Leftrightarrow \forall x \in X \cdot$

```
15 terminated(x,S2) :- executed(a,S1), action(a,d),
16                     state(S1), state(S2), S2=S1+d-1.
```

Creating the goals: $\forall g \in G$

```
17 goal(g,v).
18 satisfied(g,S) :- EX(g^+,S), not EX(g^-,S), state(S).
```

Creating the norms: $\forall n = \langle o | f, a_{sub}, a_{con}, dl, c \rangle \in N$

```
19 norm(n,c).
```

$\forall n = \langle o, a_{sub}, a_{con}, dl, c \rangle \in N$

```
20 holdsat(o(n,S1,a_sub,dl+S2),S2) :- executed(a_con,S1),
21              action(a_con,d), S2=S1+d,state(S1), state(S2).
22 cmp(o(n,S1,a,DL),S2) :- holdsat(o(n,S1,a,DL),S2),
23    executed(a,S2),action(a,d),state(S1),state(S2),S2!=DL.
24 terminated(o(n,S1,a,DL),S2) :- cmp(o(n,S1,a,DL),S2),
25                                state(S1), state(S2).
26 vol(o(n,S1,a,DL),S2) :- holdsat(o(n,S1,a,DL),S2), DL=S2,
27                         state(S1), state(S2).
28 terminated(o(n,S1,a,DL),S2) :- vol(o(n,S1,a,DL),S2),
29                                 state(S1), state(S2).
```

$\forall n = \langle f, a_{sub}, a_{con}, dl, c \rangle \in N$

```
30 holdsat(f(n,S1,a_sub,dl+S2),S2) :- executed(a_con,S1),
31               action(a_con,d),S2=S1+d,state(S1), state(S2).
32 cmp(f(n,S1,a,DL),S2) :- holdsat(f(n,S1,a,DL),S2),
33                  action(a,d), DL=S2, state(S1), state(S2).
34 terminated(f(n,S1,a,DL),S2) :- cmp(f(n,S1,a,DL),S2),
35                                state(S1), state(S2).
36 vol(f(n,S1,a,DL),S) :- holdsat(f(n,S1,a,DL),S2),
37              executed(a,S2),state(S1) state(S2), S2!=DL.
38 terminated(f(n,S1,a,DL),S2) :- vol(f(n,S1,a,DL),S2),
39                                state(S1), state(S2).
```

Plans need to satisfy at least one goal

```
40 satisfied(g) :- satisfied(g,S), state(S).
41 :- not satisfied(g1), ... , not satisfied(gm).
```

Avoiding conflicting actions: $\forall\ (a_1, a_2) \in cf_{action}$

```
42 :- inprog(a1,S),inprog(a2,S),action(a1,d1), action(a2,d2),
43         state(S).
```

Figure 2: Mapping $P = (FL, I, A, G, N)$ to its Corresponding Computational Model Π_{PBase}

execution that is the subject of the norm, or at the end state of action's execution. In the encoding we show an implementation of the former; the latter can be catered for in a similar fashion.

Lines 19–39 deal with obligations and prohibitions of the form $n = \langle d_o, a_{con}, a_{sub}, dl, c \rangle$. Line 19 encodes norm n with penalty cost c upon violation. In order to implement the concepts of norm compliance and violation for instantiated norms, we introduce a normative fluent that holds over the compliance period. The compliance period begins from the state in which action a_{con}'s execution ends. The compliance period then ends within dl time units of the end of action a_{con}, which is denoted as dl' in the normative fluent. For instance, fluent $o(n_1, s', a_{sub}, dl')$ expresses that the instance of norm n_1 that was activated in state s', obliges the agent to execute action a_{sub} before deadline dl'. The state in which the norm is activated is a part of the fluent to distinguish different activations of the same norm from one another. For example, fluent $o(n_1, s'', a_{sub}, dl'')$ refers to a different instance of norm $n1$ that was activated in s''. An obligation fluent denotes that action a_{sub}'s execution should begin before deadline dl' or be subject to violation, while prohibition fluent $f(n_2, s', a_{sub}, dl')$ denotes that action a_{sub} should not begin before deadline dl' or be subject to violation. Lines 20–21 and 30–31 establish respectively the obligation and prohibition fluents that hold for the duration of the compliance period.

In terms of compliance, if the obliged action *starts* during the compliance period in which the obligation fluent holds, the obligation is complied with (line 22–23). Compliance is denoted by the atom `cmp`. The obligation fluent is terminated in the same state that compliance is detected (lines 24–25). If the deadline expires and the obligation fluent still holds, it means that the compliance never occurred during the compliance period and the norm is therefore violated (lines 26–27). The atom `vol` denotes violation. The obligation fluent is terminated when the deadline expires and the norm is violated (lines 28–29).

On the other hand, a prohibition norm is violated if the forbidden action *begins* during the compliance period in which the prohibition fluent holds (lines 36–37). As with the obligation norms, after being violated, the prohibition fluent is terminated (lines 38–39). If the deadline expires and the prohibition fluent still holds, that means the prohibited action did not begin during the compliance period and the norm is therefore complied with (lines 32–33). The obligation fluent is terminated in the same state that compliance is detected (lines 34–35).

4.2.4. Mapping Answer Sets to Plans

Having implemented the components of $P = (FL, \Delta, A, G, N)$, we now encode the criteria for a sequence of actions to be identified as a plan and a solution to P. The rule in line 41 is responsible for constraining the answer sets to those that fulfill at least one goal. This is done by excluding answers that do not satisfy any goal. The input for this rule is provided in line 40, where goals are marked as satisfied if they are satisfied in at least one state. Prevention of concurrent conflicting actions is achieved via lines 42–43 which establish that two such actions cannot be in progress simultaneously. This concludes the mapping of a formal planning problem to its computational counterpart in *AnsProlog*. For a problem P we refer to the program consisting of lines 1–43 as Π_{PBase}.

As mentioned in Section 2, the formulation of our Disaster scenario is provided in Appendix A. The mapping of the scenario to its computational model follows in Appendix B.

4.2.5. Soundness and Completeness of Implementation

The following theorems state the correspondence between the solutions for problem P and answer sets of program Π_{PBase}.

Theorem 1 (Soundness)**.** *Let $P = (FL, I, A, G, N)$ be a normative practical reasoning problem with Π_{PBase} as its corresponding AnsProlog program. Let Ans be an answer set of Π_{PBase}, then a set of atoms of the form $executed(a_i, t_{a_i}) \in Ans$ encodes a solution to P.*

The proof of this theorem is presented in Appendix C. This is a proof by structure that explains how the structure of Π_{PBase} satisfies the conditions that identifies a sequence of actions as a plan.

Theorem 2 (Completeness)**.** *Let $\pi = \langle (a_0, 0), \cdots, (a_n, t_{a_n}) \rangle$ be a plan for $P = (FL, I, A, G, N)$. Then there exists an answer set of Π_{PBase} containing atoms $executed(a_i, t_{a_i})$ that correspond to π.*

The proof of this theorem is presented in Appendix D. In this proof the program is first transformed to a program without any naf-literals and choice rules. We then take a candidate answer set for the program and show that it is a minimal model for the transformed program.

```
44 value(TV) :- TV = #sum {V: goal(G,V), satisfied(G)}.
45 violated(N,S1) :- vol(o(N,S1,a,DL),S2), state(S1;S2).
46 violated(N,S1) :- vol(f(N,S1,a,DL),S2), state(S1;S2).
47 cost(TC) :- TC = #sum{C,S:violated(N,S),norm(N,C)}.
48 utility(TV-TC) :- value(TV), cost(TC).
49 #maximize {U:utility(U)}.
```

Figure 3: Optimal Plans for $P = (FL, I, A, G, N)$

4.2.6. Optimal Plans

In order to find optimal plans in Figure 3 we show how to encode the utility function defined by Eq. 19. The sum of values of goals satisfied in a plan is calculated in line 44, where we use an ASP built-in aggregate #`sum`. This aggregate is an operation on a set of weighted literals that evaluates to the sum of the weights of the literals. We first assign the value of goals as the weight of literals `satisfied(G)` and then use #`sum` to compute the sum of value of all goals satisfied.

The sum of costs of norms violated in a plan is calculated in line 47 using the same aggregate. However, the weight of the literal is the cost of punishment of the norms. The input for this line is given in lines 45 and 46, where violated norms are distinguished based on the combination of the norm id n and the state s in which they are instantiated. Having calculated `value(TV)` and `cost(TC)`, the utility of a plan is computed in line 48, which is subject to a built-in optimisation statement in the final line. This optimisation statement identifies an answer set as optimal if the sum of weights of literals that hold is maximal with respect to all answer sets of the program. By assigning `U` as the weight of literal `utility(U)` we compute the answer set that maximises the utility.

Let program $\Pi_P = \Pi_{PBase} \cup \Pi^*_P$, where Π^*_P consists of lines 44–49. The following theorem states the correspondence between the plans for problem P and answer sets of program Π_P.

Theorem 3. *Given a normative practical reasoning problem $P = (FL, I, A, G, N)$, for each answer set Ans of Π the set of atoms of the form executed(a_i, t_{a_i}) in Ans encodes an optimal solution to P. Conversely, each solution to the problem P corresponds to a single answer set of π.*

This theorem follows immediately from Theorem 1 and 2 and the structure of program Π^*_P.

5. Related Work

In this section we first review a number of architectures aimed at incorporating normative reasoning into practical reasoning which we group into BDI (Section 5.1) and non-BDI (Section 5.2). Then in Section 5.3 we examine different classifications of these architectures according to the approaches taken. In the same section we also compare these approaches with the approach proposed here.

5.1. BDI Architectures with Normative Reasoning

There is a substantial body of work on the integration of norms into the BDI architecture (Rao and Georgeff, 1995), but motivation, theory and practice vary substantially. A key assumption here is that a plan library (i.e. a set of pre-defined plans) exists and the agent uses normative considerations to choose and/or customise a provided plan, rather than generating a norm-compliant plan. We review several normative BDI frameworks and compare them with the approach proposed here.

The BOID architecture (Broersen et al., 2001) extends BDI with the concept of obligation and uses agent types such as social, selfish, etc. to handle the conflicts between beliefs, desires, intentions and obligations. For example, selfish agents give priority to their desires in case of conflict, whereas social agents give priority to their obligations. Since beliefs, desires, intentions and obligations are all represented as a set of rules, priorities are assigned to rules and subsequently used to resolve the conflict. BOID's rich rule-based language makes for a very expressive system, but it is now of largely historical interest, since the implementation is complicated – no reference version is currently available – and has a high computational complexity.

NoA (Kollingbaum, 2005) is a normative language and agent architecture. As a language it specifies the normative concepts of obligation, prohibition and permission to regulate a specific type of agent interaction called "supervised interaction". As a practical reasoning agent architecture, it describes how agents select a plan from a pre-generated plan library such that the norms imposed on the agent at each point of time are observed. NoA agents do not have internal motivations such as goals or values that might conflict with norms, therefore the agent will always be norm compliant. Publications on NoA do not discuss its evaluation and validate the implementation using examples, which makes the status of the implementation unclear.

ν-BDI (Meneguzzi et al., 2015) enables BDI agents to perform normative reasoning for the purpose of customising pre-existing plans that ensure compliance with the set of norms imposed on the agent. That said, there are mechanisms in place to allow norm violation where goal achievement would not otherwise be possible. In contrast to BOID and NoA, much attention is paid to the practicality and computational efficiency of reasoning in ν-BDI. This is evidenced by the complexity analysis of the algorithms for the norm management mechanism, complemented by empirical analysis, and which both report reasonable (sic) computational costs.

N-2APL (Alechina et al., 2012) is a norm-aware agent programming language based on 2APL (Dastani, 2008) that supports representation of and reasoning about beliefs, goals, plans, norms, sanctions and deadlines. The N-2APL agents select plans to execute such that they fulfill obligations imposed on agents. The agent can also choose to suppress certain plans to avoid violating prohibitions. Scheduling of plans is conducted based on plan deadlines or possible sanctions associated with the plans. The scheduling does not concern itself with construction of interleaving plans, thus scheduling boils down to sequencing of plans. Norms in N-2APL are quite simple, being either obligations to carry out and prohibitions not to carry out a specified action. The norms are not conditional (i.e. there is no activation condition defined that triggers the norm) and unlike obligations the prohibitions do not have deadlines.

N-Jason (Lee et al., 2014) sets out an extension of the Jason (Bordini et al., 2007) variant of the BDI architecture to account for norms in plan selection and to handle priorities in plan scheduling. Like N-2APL, N-Jason enables the underlying implementation of the BDI architecture to carry out norm-aware deliberation and provides a run-time norm execution mechanism, through which new unknown norms – as long as they pertain to known actions – are recognized and can trigger plans. To be able to process a norm such as an obligation, which includes a deadline, the agent architecture must be able to reason about deadlines and priorities, and choose between plans triggered by a particular norm. Consequently, N-Jason extends the syntax of the plan library to allow priority annotation, and the scheduling algorithm of AgentSpeak(RT) to operate in the context of Jason/AgentSpeak(L) and provide "real-time agency".

5.2. Non-BDI Architectures with Normative Reasoning

A second smaller group of research work considers the problem either from a non-BDI perspective or are agent-architecture agnostic. One result of not being tied to BDI is not necessarily relying on a pre-generated plan library, which raises issues of how to generate norm-compliant plans.

Oren et al. (2011) do utilise a pre-generated plan, like the list above, but take norms into consideration when deciding how to execute the plan with respect to the norms triggered by that plan. Specifically, the approach aims to adjust the chosen plan to account for the norms that govern the plan actions at each point in time, where the norm expresses constraints on the values that can be assigned to variables in a plan action. The adjustments of values in actions specify how the agent should execute a plan, such that the cost of violated norms is outweighed by the reward from norm compliance. The most preferred plan is therefore the one that maximises this metric.

Panagiotidi et al. (2012a) take norms into account in plan generation where the planning problem is specified in PDDL 2.1 (Fox and Long, 2003). The normative state of the agent is checked, using the planning tool Metric-FF (Hoffmann and Nebel, 2001), after each individual action; then the planner decides if the agent should comply with a norm or not based on a linear cost function specified in terms of constraints on the states achieved during each plan. Although this mechanism enables an agent to cope with the dynamics of operating in an open environment, checking the normative position of an agent after each action imposes a high computational cost on the plan generation phase.

Shams et al. (2015) define an approach to practical reasoning that considers norms in both plan generation and plan selection. The agent attempts to satisfy a set of potentially conflicting goals in the presence of norms, as opposed to conventional planning problems that generate plans for a single goal. The main contributions of Shams et al. (2015) are (i) the introduction of an enforcement approach that is a combination of utility-based and pressure-based compliance methods (López y López et al., 2005), and (ii) formalising the conflicts between goals, between norms, and between goals and norms. There is a penalty cost for norm violation regardless of the existence of conflict. Whenever a norm is triggered, outcomes of norm compliance and violation, and their impacts on the hinderance or facilitation of other goals and norms, are both generated and compared by utility. In those cases where there are no conflicts and no goals or norms are hindered by the punishment of violation, the loss of utility drives the agent towards compliance. Plans

are selected based on the comparison of the utility of the goals satisfied and cost of norms violated in the entire plan.

5.3. Normative Reasoning Mechanisms

We now consider the approaches accounting for norms in practical reasoning in order to uncover similarities across architectures. Not surprisingly, much work stems from planning and how to take account of norms in the plan identification or construction process. There are broadly three approaches:

1. Choosing a plan that is norm compliant (e.g., NoA (Kollingbaum, 2005)), which is a one-off process, that may fail delivering the best (where "best" can be defined in various ways) plan available for the situation from those available, and which requires starting again when a plan step fails and the remainder of the plan is invalidated. The main points of difference between NoA and the work presented here are that (i) NoA agents are BDI specific, (ii) they do not have internal motivations such as goals or values that might conflict with norms, which therefore enables the NoA agent to always comply with norms (iii) plans are pre-existing rather than generated.
2. Customising a plan to make it norm compliant (e.g., (Oren et al., 2011)) is potentially more flexible in making use of available plans (also helping customize existing plans into optimal norm-compliant plans), but otherwise has the same replanning drawback. In common with Oren et al. (2011), we use the utility of the entire plan in the selection process, but differ in that we generate plans rather than use plans from a library.
3. Generating a plan that is norm compliant (e.g., (Panagiotidi et al., 2012a; Shams et al., 2015)). The former addresses on-going compliance and re-planning, putting a high computational overhead on each plan step. Of necessity, Panagiotidi et al. (2012a) can only compute utility on a step-by-step basis, whereas we consider the utility of the whole plan. Shams et al. (2015) attempt to balance compliance on the part of the agent (where the agent chooses a norm-compliant action in preference) with enforcement (where the agent is discouraged from non-norm-compliance via punishments for norm violation), but is not robust to plan failure. Furthermore, in Shams et al. (2015), conflict is formulated in advance by taking a static view about conflicts. For instance, two goals that are logically inconsistent, cannot be satisfied in the same plan, regardless of the order or choice of actions in a plan.

Table 2: Summary of Related Frameworks

Framework	Deontic Operator	Activation Condition	De-activation Condition
BOID (Broersen et al., 2001)	o	N/A	N/A
NoA (Kollingbaum, 2005)	o, f, p	state, action	state, action
ν-BDI (Meneguzzi et al., 2015)	o, f	state	state
N-2APL (Alechina et al., 2012)	o, f	state	state[3]
N-Jason (Lee et al., 2014)	o, p, w[4]	N/A	temporal constraint
Oren et al. (2011)	o, f	N/A	N/A
Panagiotidi et al. (2012a)	o, f	state	state
Shams et al. (2015)	o, f	action	temporal constraint
This work	o, f	action	temporal constraint

In contrast, in the work presented here, conflicts are not formulated in advance; instead, they are inferred from plans. Therefore, the agent might be able to schedule its actions such that two goals that are logically inconsistent are satisfied in the same plan at different points in time. As discussed in Section 3.6, the norm representation is extended to accommodate compliance and violation in the presence of durative actions more flexibly by allowing compliance to be defined as the start or end of the action that is the subject of the norm.

Table 2 shows a summary of related framework to the framework proposed in this paper. For the current work, the parameters compared are the same as Shams et al. (2015). The majority of frameworks, including our framework, deal with obligations and prohibitions. Activation condition, however, varies in those that do support conditional norms. An activation condition presented as an action can be expressed as a state that satisfies the post-conditions of the action. Unlike many related frameworks, we have exploited the explicit representation of time in formal model to encode the norm de-activation condition as a time instant. As discussed in Section 3.5, associating a deadline with temporal properties is considered to be realistic and dynamic, in particular when the norms capture the requirements of real-world scenarios (Chesani et al., 2013; Kafali et al., 2014; Gasparini et al., 2015), such as the disaster scenario we have modelled in this paper. Another important differentiation point between our work and the related ones

[3]The de-activation condition only applies to obligations. Prohibitions do not have such a condition.

[4]Operator w stands for power and indicates the capability of doing something in prohibitive societies, where actions are not allowed unless empowered and explicitly permitted.

is that our model is capable of handling durative actions and their execution concurrently, as well as dealing with norm compliance and violation in the presence of durative actions.

6. Conclusions and Future Work

An agent performing practical reasoning in an environment regulated by norms constantly needs to weigh up the importance of goals to be satisfied and norms with which to comply. This decision process is only possible when the agent has access to the set of all possible plans available and the agent can ascertain the impact of its decision on entire plans. This research offers means to capture and measure the impact via utility functions, offering numeric metrics, so that the decision problem can be reformulated as choosing a plan from a set of generated plans, which maximises its overall utility. While the literature we have surveyed contains practical reasoning frameworks that take into account normative considerations, they are limited in several ways, and we have contrasted them with our approach in this article.

The majority of these frameworks are limited to a specific type of agent architecture, mostly BDI, (e.g., (Broersen et al., 2001; Kollingbaum, 2005; Meneguzzi et al., 2015)). In our research we do not assume any specific architecture. Instead, we adopt a realistic view that agents have capabilities (encoded as the actions they may perform), and they have internal motivations (encoded as the goals of a planning problem). This leaves the option of extending current work to a multi-agent setting where agents might not necessarily have the same architecture.

In the approaches set out in the literature the attitude agents have towards norms is often one of compliance, meaning that their plans are often selected or, in some approaches, customised, to ensure norm compliance, (e.g., (Kollingbaum, 2005; Alechina et al., 2012; Oren et al., 2011)). We argue that in certain situations an agent might be better off violating a norm which, if followed, would make it impossible for the agent to achieve an important goal or complying with a more important norm; we enable agents to compare the relative importance of goals and norms via their utilities. Consequently, through our practical reasoning process agents consider *all* plans (i.e., sequences of actions), including those leading to norm compliance and violation; each plan gets an associated overall utility for its sequences of actions, and norms followed/violated, and agents can decide which of them to pursue. The plan an agent chooses to follow is not necessarily norm-

compliant, however, our mechanism guarantees that the decision will maximise the overall plan/norm utility, which justifies the occasional violation of norms as the plan is followed.

We see several interesting avenues for future work. Our research currently addresses normative practical reasoning in a single-agent setting, extending to a multi-agent setting seems a natural next step. This idea can be explored both when the agents are collaborating to fulfill common goals, as well as when they are competing to use resources to fulfill their individual goals. In the former case, the best course of action can be identified as one that maximises the overall utility of the system. In the latter, game-theoretic approaches can be utilised to identify a solution that ensure relative optimality for the individuals (e.g. Ågotnes et al. (2007)). Another possibility to explore in a multi-agent setting is to infer conflicts between goals, between norms and between goals and norms by analysing the overall set of possible plans. The inferred conflicts can guide the process of re-engineering of the system toward a more social and norm compliant system (e.g. (Savarimuthu et al., 2013)).

We note the relative limitations of our norm representation. Although our approach addressed action-based norms, we envisage how it can be extended and adapted to handle state-based norms. Our Def. 4 needs to cater for formulae to represent both the norm activation condition, a_{con}, and the norm subject, a_{sub}, instead of actions. A combination of action- and state-based norms (e.g. De Vos et al. (2013)) enriches the norm representation as well as normative reasoning. Also, the norm representation language can be extended to cater for deadlines that are expressed as reaching a state[5] rather than a time instance. For instance, an obligation to open a dam on a river can come in force when the water level is above a certain point, and subsequently terminated when the water level drops below a certain level, regardless of how long it takes for that to happen. We would also like to include permission norms in addition to obligations and prohibitions. The modelling of permissions as exceptions to obligations and prohibitions has been used to justify violations under specific circumstances, (e.g. Oren et al. (2010); Criado (2012)). It is also used as a means to handle the uncertainty and incompleteness of the knowledge of the environment the agents operate in (Alrawagfeh and Meneguzzi, 2014).

[5]In such cases the deadline is referred to as a norm termination condition.

Finally, our implementation should be seen as a proof-of-concept that, apart from replanning, provides a provable computational realisation of all aspects of the formal model. In future, we aim at extending the implementation to accommodate replanning when a plan in progress is interrupted for any reason. The formal model is implementation language neutral so other implementation languages could be used.

Appendix A. Formulation of the Disaster Scenario

We provide a formalisation of the scenario set out in Section 2. Let $P = (FL, \Delta, A, G, N)$ be the normative practical reasoning problem for the disaster scenario such that:

- $FL = \left\{ \begin{array}{c} shockDetected, poisonDetected, waterSupplied, \\ areaSecured, evacuated, shockDetected, \\ shelterBuilt, populated, wounded, \\ earthquakeDetected, medicineSupplied, \\ noAccess, medicsPresent \end{array} \right\}$
- $\Delta = \left\{ \begin{array}{c} earthquakeDetected, medicsPresent, \\ wounded, populated, waterSupplied \end{array} \right\}$
- $A = \left\{ \begin{array}{c} detectShock, detectPoison, stopWater, \\ buildShelter, evacuate, getMedicine, secure \end{array} \right\}$ where

  $detectShock = \langle \{\}, \{shockDetected\}, 1 \rangle$.

  $detectPoison = \langle \{\}, \{poisonDetected\}, 1 \rangle$.

  $stopWater = \left\langle \left\{ \begin{array}{c} poisonDetected, \\ waterSupplied \end{array} \right\}, \{\neg waterSupplied\}, 1 \right\rangle$

  $buildShelter = \left\langle \left\{ \begin{array}{c} areaSecured, \\ evacuated, \\ \neg shockDetected \end{array} \right\}, \left\{ \begin{array}{c} shelterBuilt, \\ \neg evacuated \end{array} \right\}, 4 \right\rangle$

  $evacuate = \left\langle \left\{ \begin{array}{c} shockDetected, \\ populated \end{array} \right\}, \left\{ \begin{array}{c} evacuated, \\ \neg populated \end{array} \right\}, 5 \right\rangle$

  $getMedicine = \left\langle \left\{ \begin{array}{c} earthquakeDetected, \\ wounded \end{array} \right\}, \{medicine\}, 3 \right\rangle$

  $secure = \left\langle \{evacuated\}, \left\{ \begin{array}{c} areaSecured, \\ noAccess \end{array} \right\}, 3 \right\rangle$

- $G = \{runningHospital, organiseSurvivorCamp\}$, where:

  $runningHospital = \left\langle \left\{ \begin{array}{c} medicsPresent, \\ waterSupplied, \\ medicineSupplied \end{array} \right\}, 25 \right\rangle$

  $organiseSurvivorCamp = \left\langle \left\{ \begin{array}{c} areaSecured, \\ shelterBuilt \end{array} \right\}, 18 \right\rangle$

- $N = \{n_1, n_2\}$, where:

  $n_1 = \langle f, detectShock, buildShelter, 3, 5 \rangle$

  $n_2 = \langle o, detectPoison, stopWater, 2, 10 \rangle$

Appendix B. Mapping of Our Disaster Scenario

The formal specification of our disaster scenario (Section2) as provided in the previous section can be mapped to its corresponding *AnsProlog* program following the rules given in Figure 2 on page 25. The corresponding program is shown in Figures B.4-B.8. Optimisation conditions are shown in Figure B.9. The visualisation of the three answer sets of the program is displayed in figures B.10–B.12, where arcs show the actions in progress and the boxes below each state, show the fluents that hold in that state. The fluents in bold are the fluents that are added to the state, while the crossed fluents are the terminated ones. Norms violated in a state are highlighted in red and goals satisfied are highlighted in green. Applying the optimisation statements in Figure B.9, the utility of each plan presented by each answer set is calculated as below, making the plan presented by answer set 3 the optimal plan.

Utility of plan presented by answer set 1, Figure B.10:
$Utility(\pi_{Ans1}) = v(runningHospital) = 25$

Utility of plan presented by answer set 2, Figure B.11:
$Utility(\pi_{Ans2}) = v(runningHospital) + v(organiseSurvivorCamp) - c(n_1) - c(n_1) = 25 + 18 - 5 - 5 = 33$

Utility of plan presented by answer set 3, Figure B.12:
$Utility(\pi_{Ans3}) = v(runningHospital) + v(organiseSurvivorCamp) = 43$

```
%Rules for states

state(0..13).
holdsat(earthquakeDetected,0).
holdsat(medicsPresent,0).
holdsat(wounded,0).
holdsat(populated,0).
holdsat(waterSupplied,0).

%Inertial fluents

holdsat(X,S2) :- holdsat(X,S1), not terminated(X,S1),
                 state(S1), state(S2), S2=S1+1.

%Rules for Translating Actions

action(detectShock,1).
pre(detectShock,S) :- state(S).
holdsat(shockDetected,S2) :-
   executed(detectShock,S1), action(detectShock,1),
   state(S1), state(S2), S2=S1+1.

action(detectPoison,1).
pre(detectPoison,S) :- state(S).
holdsat(poisonDetected,S2) :-
        executed(detectPoison,S1),
        action(detectPoison,1), state(S1),
        state(S2), S2=S1+1.

action(stopWater,1).
pre(stopWater,S) :- holdsat(poisonDetected,S),
            holdsat(waterSupplied,S),state(S).
terminated(waterSupplied,S2) :-
       executed(stopWater,S1), action(stopWater,1),
       state(S1), state(S2), S2=S1.
```

Figure B.4: The Corresponding *AnsProlog* of The Disaster Scenario Part 1

```
36 action(buildShelter,4).
37 pre(buildShelter,S) :- holdsat(areaSecured,S),
38                holdsat(evacuated,S), state(S).
39 holdsat(shelterBuilt,S2) :-
40  executed(buildShelter,S1), action(buildShelter,4),
41  state(S1), state(S2), S2=S1+4.
42 terminated(evacuated,S2) :-
43  executed(buildShelter,S1), action(buildShelter,4),
44  state(S1), state(S2),S2=S1+3.
45
46 action(evacuate,5).
47 pre(evacuate,S) :- holdsat(populated,S),
48                    holdsat(shockDetected,S),state(S).
49 holdsat(evacuated,S2) :- executed(evacuate,S1),
50         action(evacuate,5), state(S1), state(S2),
51         S2=S1+5.
52 terminated(populated,S2) :- executed(evacuate,S1),
53         action(evacuate,5), state(S1), state(S2),
54         S2=S1+4.
55
56 action(getMedicine,3).
57 pre(getMedicine,S) :-
58                    holdsat(earthquakeDetected,S),
59                    holdsat(wounded,S),state(S).
60 holdsat(medicineSupplied,S2) :-
61                   executed(getMedicine,S1),
62                   action(getMedicine,3),
63                   state(S1), state(S2), S2=S1+3.
64
65 action(secure,3).
66 pre(secure,S) :- holdsat(evacuated,S),state(S).
67 holdsat(areaSecured,S2) :- executed(secure,S1),
68         action(secure,3), state(S1), state(S2),
69         S2=S1+3.
70 holdsat(noAccess,S2) :- executed(secure,S1),
71       action(secure,3), state(S1), state(S2),
72       S2=S1+3.
```

Figure B.5: The Corresponding *AnsProlog* of The Disaster Scenario Part 2

```
{executed(A,S)} :- action(A,D), state(S).
inprog(A,S2) :- executed(A,S1), action(A,D),
            state(S1), state(S2), S1 <= S2, S2< S1+D.
:- inprog(A,S), action(A,D), not pre(A,S), state(S).

% This rule ensures that not more than one action can
% be executed per state.
:- executed(A1,S), executed(A2,S), A1!=A2,
   action(A1,D1), action(A2,D2), state(S).

% These rules ensure that there is no action executed
% or is in progress in the final state.
:- executed(A,13), action(A,D).
:- inprog(A,13), action(A,D).

% This rule ensures that the agent is not
% idle at any point.
alpha(S) :- inprog(A,S), action(A,D), state(S).
:-  not alpha(S), state(S), S<13.

%Rules for Translating Norms

norm(n1,5).
norm(n2,10).

holdsat(f(n1,S1,buildShelter,3+S2),S2) :-
     executed(detectShock,S1),
     action(detectShock,1),
     S2=S1+1, state(S1), state(S2).
cmp(f(n1,S1,buildShelter,DL),S2) :-
     holdsat(f(n1,S1,buildShelter,DL),S2),
     action(buildShelter,4), DL=S2,
     state(S1), state(S2).
terminated(f(n1,S1,buildShelter,DL),S2) :-
     cmp(f(n1,S1,buildShelter,DL),S2),
     state(S1), state(S2).
```

Figure B.6: The Corresponding *AnsProlog* of The Disaster Scenario Part 3

```
109 vol(f(n1,S1,buildShelter,DL),S2) :-
110      holdsat(f(n1,S1,buildShelter,DL),S2),
111      executed(buildShelter,S2),state(S1),
112      state(S2), S2!=DL.
113 terminated(f(n1,S1,buildShelter,DL),S2) :-
114      vol(f(n1,S1,buildShelter,DL),S2),
115      state(S1), state(S2).
116
117 holdsat(o(n2,S1,stopWater,2+S2),S2) :-
118      executed(detectPoison,S1),
119      action(detectPoison,1), S2=S1+1,
120      state(S1), state(S2).
121 cmp(o(n2,S1,stopWater,DL),S2) :-
122      holdsat(o(n2,S1,stopWater,DL),S2),
123      executed(stopWater,S2),
124      action(stopWater,1),state(S1),
125      state(S2), S2!=DL.
126 terminated(o(n2,S1,stopWater,DL),S2) :-
127      cmp(o(n2,S1,stopWater,DL),S2), state(S1),
128      state(S2).
129 vol(o(n2,S1,stopWater,DL),S2) :-
130      holdsat(o(n2,S1,stopWater,DL),S2), DL=S2,
131      state(S1), state(S2).
132 terminated(o(n2,S1,stopWater,DL),S2) :-
133      vol(o(n2,S1,stopWater,DL),S2), state(S1),
134      state(S2).
135
136 % Generating Solution for problem P
137
138 goal(runningHospital,25).
139 satisfied(runningHospital,S) :-
140  holdsat(medicsPresent,S),holdsat(waterSupplied,S),
141  holdsat(medicineSupplied,S), state(S).
142 satisfied(runningHospital) :-
143              satisfied(runningHospital,S), state(S).
```

Figure B.7: The Corresponding *AnsProlog* of The Disaster Scenario Part 4

```
goal(organiseSurvivorCamp,18).
satisfied(organiseSurvivorCamp,S) :-
                  holdsat(areaSecured,S),
                  holdsat(shelterBuilt,S), state(S).
satisfied(organiseSurvivorCamp) :-
     satisfied(organiseSurvivorCamp,S), state(S).

:- not satisfied(runningHospital),
   not satisfied(organiseSurvivorCamp).

:- inprog(detectShock,S), inprog(buildShelter,S),
   action(detectShock,1), action(buildShelter,4),
   state(S).

:- inprog(buildShelter,S),  inprog(evacuate,S),
   action(buildShelter,4), action(evacuate,5),
   state(S).

:- inprog(buildShelter,S),  inprog(secure,S),
  action(buildShelter,4), action(secure,3), state(S).
```

Figure B.8: The Corresponding *AnsProlog* of The Disaster Scenario Part 5

```
violated(n1,S1) :- vol(f(n1,S1,buildShelter,DL),S2),
                   state(S1), state(S2).
violated(n2,S1) :- vol(o(n2,S1,stopWater,DL),S2),
                   state(S1), state(S2).

%Optimisation rules

value(TV) :- TV = #sum{V: goal(G,V),satisfied(G)}.
cost(TC) :- TC = #sum{C,S:violated(N,S),norm(N,C)}.

utility(TV-TC) :- value(TV),cost(TC).

#maximize{U:utility(U)}.
```

Figure B.9: Optimisation Rules for The Disaster Scenario

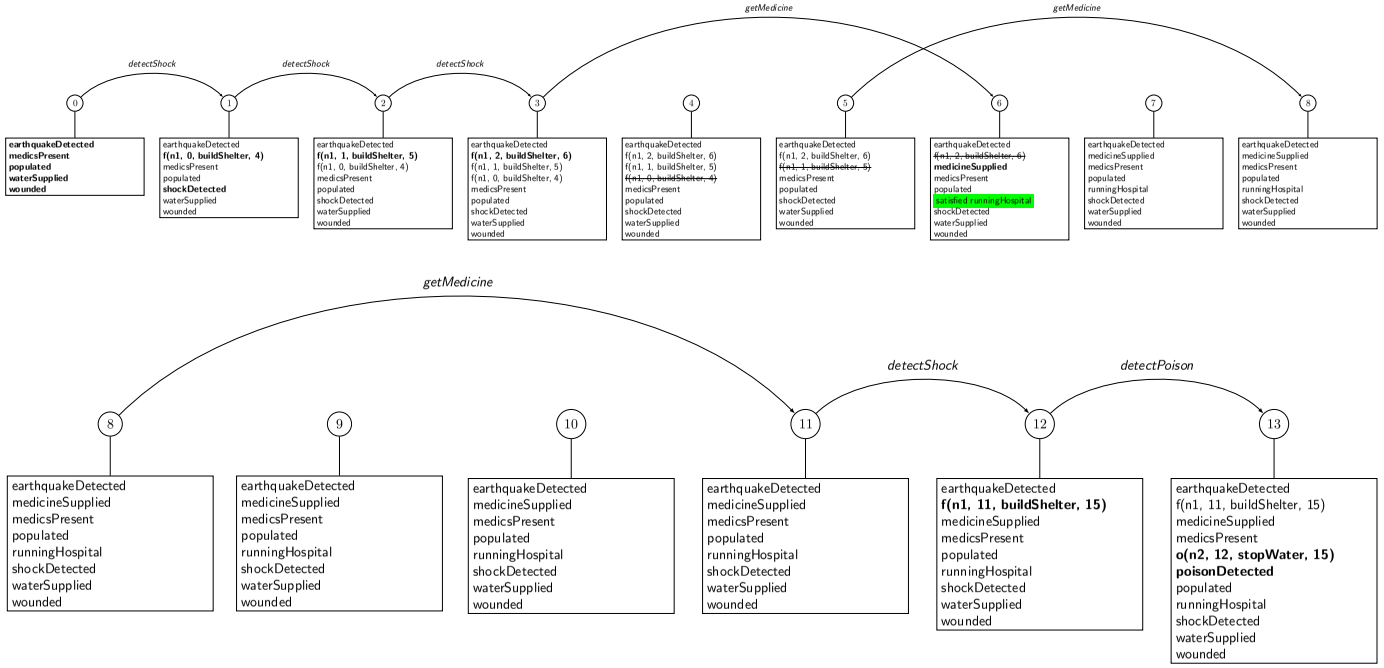

Figure B.10: Visualisation of answer set 1

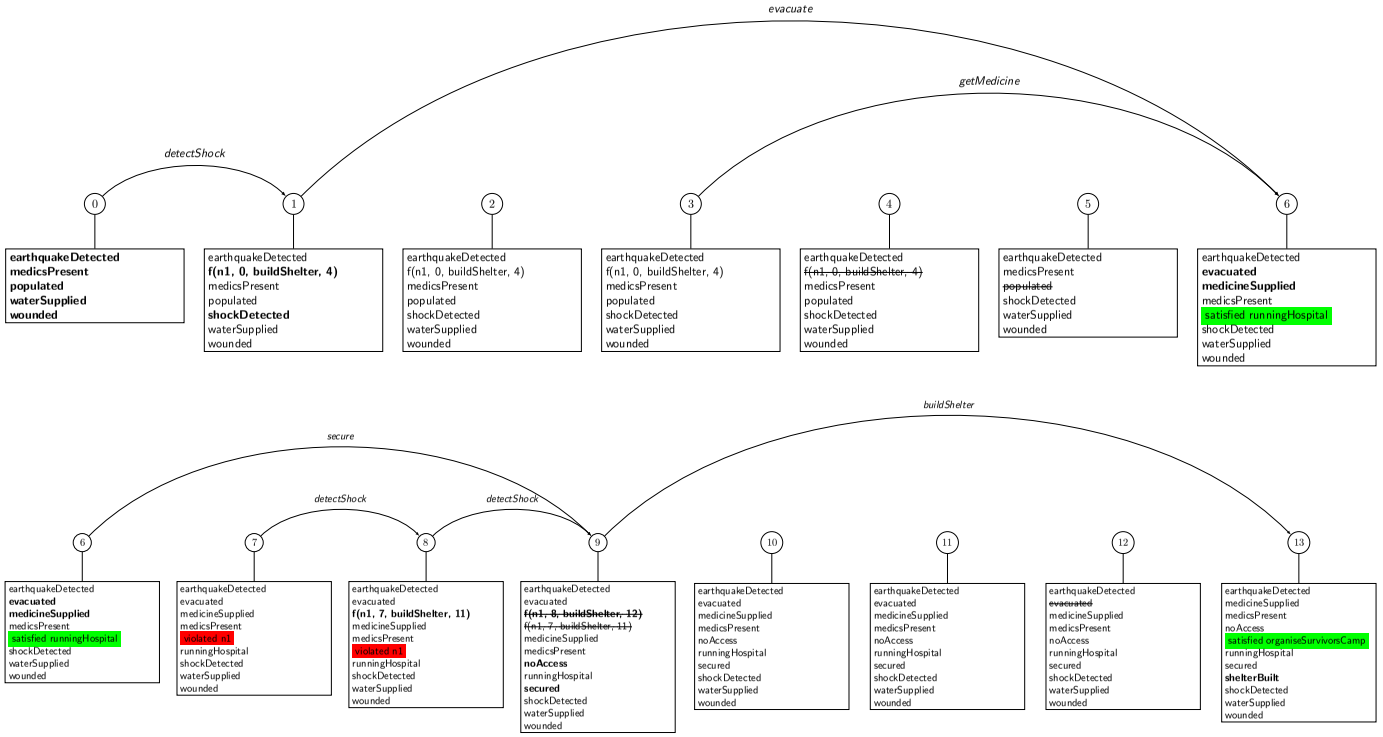

Figure B.11: Visualisation of answer set 2

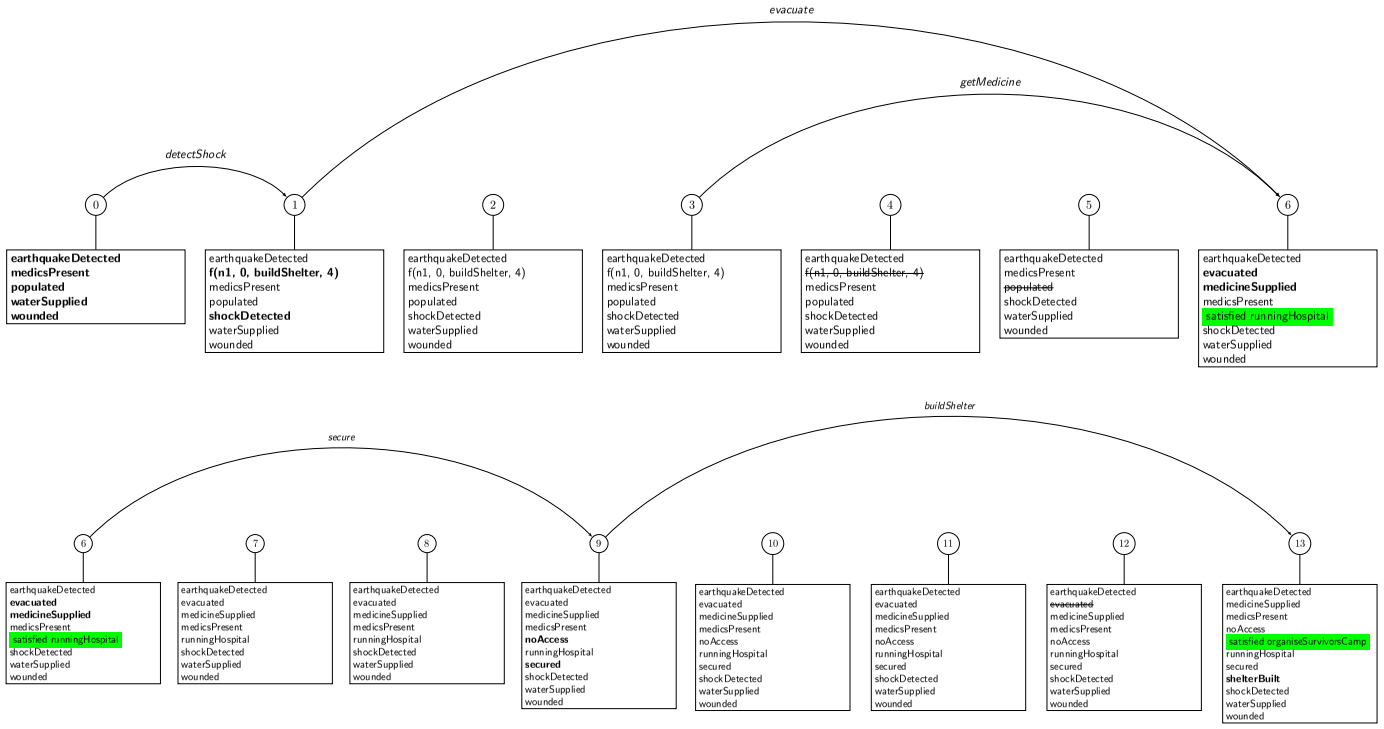

Figure B.12: Visualisation of answer set 3

Appendix C. Proof of Theorem 1

Theorem (Soundness)**.** *Let $P = (FL, I, A, G, N)$ be a normative practical reasoning problem with Π_{PBase} as its corresponding AnsProlog program. Let Ans be an answer set of Π_{PBase}, then a set of atoms of the form executed(a_i, t_{a_i}) $\in$ Ans encodes a solution to P.*

Proof. We need to prove that program Π_{PBase} (Figure 2) generates all sequences of actions that meet the criteria that identifies a sequence of actions as a plan, as defined in Definition 15. This implies that the sequence of actions that is a part of the answer set satisfies all the criteria to be a solution to the encoded planning program.

Actions and more precisely the postconditions of actions are what cause the change from one state to another one. Line 7 generates all sequences of actions. Lines 13–14 changes a state in which some actions end by adding the add postconditions of those actions to the state. In contrast, Lines 15 and 16 terminate the delete postconditions of actions ending in the next state such that those postconditions do not hold in the following state. If there is no action ending in a state, the state remains unchanged as all the fluents are inertial and they hold in the next state unless they are terminated (Line 4). A sequence of actions $\pi = \langle (a_0, 0), \ldots, (a_n, t_{a_n})\rangle$ s.t. $\nexists$ $(a_i, t_{a_i}), (a_j, t_{a_j}) \,\widehat{\in}\, \pi$ s.t. $t_{a_i} \leq t_{a_j} < t_{a_i} + d_{a_i}, (a_i, a_j) \in cf_{action}$ is a plan for the normative planning problem $P = (FL, \Delta, A, G, N)$ if the following conditions hold:

- fluents in Δ (and only those fluents) hold in the initial state: $s_0 = \Delta$.
  Line 2 ensures that all fluents in Δ are added to the initial state s_0.
- the preconditions of action a_i hold at time t_{a_i} and throughout the execution of a_i: $\forall k \in [t_{a_i}, t_{a_i} + d_{a_i}), s_k \models pr(a_i)$.
  Lines 10 guarantees that the preconditions of an action hold all through its execution.
- plan π satisfies a non-empty subset of goals: $G_\pi \neq \emptyset$.
  Line 41 indicates that a non-empty subset of goals has to be satisfied in a plan.

□

Appendix D. Proof of Theorem 2

Theorem (Completeness)**.** *Let $\pi = \langle (a_0, 0), \cdots, (a_n, t_{a_n})\rangle$ be a plan for $P = (FL, I, A, G, N)$. Then, there exists an answer set of Π_{PBase} containing atoms*

executed(a_i, t_{a_i}) *that corresponds to* π.

For the sake of making this proof self-contained, we first provide a formal definition of an answer set and explain the concept of reduct. As mentioned previously, a number of syntactic language representations for ASP exist. We use AnsProlog which is one of the most common classes of these languages. and it has the following elements (Baral, 2003):

Term: A term is a constant or a variable or a n-ary function $f(t_1, \cdots, t_n)$, where f is the function symbol and $t_1, \cdots, t_n$ are terms. Constants start with a lower-case letter, whereas variables start with an upper-case level. A term is *ground* if no variable occurs in it.

Atom: Atoms are the basic components of the language that can be assigned a truth value as true or false. An atom is a statement of form $A(t_1, \cdots, t_n)$, where A is a predicate symbol and $t_1, \cdots, t_n$ are terms.

Literal: Literals are atoms or negated atoms. Atoms are negated using *negation as failure* (*not*). *not a* is true if there is no evidence proving the truth of a. An atom preceded by *not* is referred to as a naf-literal.

The Herbrand universe of language $\mathcal{L}$ denoted as $HU_{\mathcal{L}}$ is the set of all ground terms which can be formed with the functions and constants in $\mathcal{L}$. The set of all ground atoms which can be formed with the functions, constants and predicates in $\mathcal{L}$ is called Herbrand base of language $\mathcal{L}$ and is denoted using $HB_{\mathcal{L}}$.

In Section 4.1, we explained that an AnsProlog program (e.g. Π) consists of a finite set of rules. In order to interpret a rule that contains variables, the rule has to be *grounded*. The grounding of each rule r in Π is then the set of all rules obtained from substitutions of elements of HU_Π for the variables in the rule r. By grounding all $r \in \Pi$, we obtain *ground*(Π).

The semantics of AnsProlog is defined in terms of *answer sets*. The answer sets of program Π, are defined in terms of the answer sets of the ground program *ground*(Π). An AnsProlog program without any naf-literal is denoted as Ansprolog^{-not}. An answer set of an AnsProlog^{-not} program Π is a minimal subset (with respect to subset ordering) S of HB that is closed under *ground*(Π). The approach to define the answer sets of an AnsProlog program Π is to take a candidate answer set S of the program and transform Π with respect to S to obtain an Ansprolog^{-not} denoted by Π^S. S is an answer set of Π if S is the answer set of AnsProlog^{-not} program Π^S. This transformation is referred to as Gelfond-Lifschitz (Gelfond and Lifschitz, 1988) transformation.

Given an AnsProlog program Π and a set S of atoms from HB_Π, the

Gelfond-Lifschitz (Gelfond and Lifschitz, 1988) transformation Π^S is obtained by deleting:

1. each rule that has a *not* L in its body with $L \in S$, and
2. literals of form *not* L in the bodies of the remaining rules.

The transformation (reduct) of choice rules was not a part of original Gelfond-Lifschitz transformation and was introduced later in (Lee et al., 2008). Recently a simplified reduct for programs including choice rules is proposed by Mark Law and Broda (2015) as follows. Given an AnsProlog program Π - with choice rules- and a set S of atoms from HB_Π, the transformation Π^S is constructed in the following 4 steps:

1. Delete each rule that has a *not* L in its body with $L \in S$.
2. Delete literals of form *not* L in the bodies of the remaining rules.
3. for any choice rule r, $l\{h_0, \cdots, h_k\}u :- body(r)$, such that $l \leq |S \cap \{h_0, \cdots, h_k\}| \leq u$, replace r with the set of rules $\{h_i :- body^+(r) | h_i \in S \cap \{h_0, \cdots, h_k\}\}$.
4. for any remaining choice rules r, $l\{h_0, \cdots, h_k\}u :- body(r)$, replace r with the constraint $:- body^+(r)$.

After these transformation, the AnsProlog program Π is a program without any naf-literals and choice rules and it is therefore, an AnsProlog^{-not}, for which the answer sets are already defined.

Proof. Let the execution of sequence of actions in $\pi = \langle (a_0, 0), \cdots, (a_n, t_{a_n}) \rangle$ bring about the sequence of states $\langle s_0, \cdots s_q \rangle$. Let M_t be the set of following atoms (and nothing else):

$$\forall k, 0 \leq k \leq q \cdot M_t \models state(k) \tag{D.1}$$

$$\forall k, 0 \leq k \leq q \cdot x \in s_k \Rightarrow M_t \models holdsat(x, k) \tag{D.2}$$

$$\forall k, 0 \leq k < q \cdot x \in (s_k \setminus s_{k+1}) \Rightarrow M_t \models terminated(x, k) \tag{D.3}$$

$$\forall a \in A \cdot M_t \models action(a, d) \tag{D.4}$$

$$\forall k, 0 \leq k \leq q \cdot a \in A, pr(a)^+ \subseteq s_k, pr(a)^- \notin s_k \Rightarrow M_t \models pre(a, k) \tag{D.5}$$

$$\forall k, 0 \leq k < q \cdot (a, k) \,\widehat{\in}\, \pi \Rightarrow M_t \models executed(a, k) \tag{D.6}$$

$$\forall a. M_t \models executed(a, k),\ \forall k, t_a \leq k < t_a + d(a) \cdot M_t \models inprog(a, k) \tag{D.7}$$

$$\forall a. M_t \models executed(a, k),\ \forall x \in ps(a)^+ \cdot M_t \models holdsat(x, k + d(a)) \tag{D.8}$$

$$\forall a.M_t \models executed(a,k),\ \forall x \in ps(a)^- \cdot M_t \models terminated(x, k + d(a) - 1) \tag{D.9}$$

$$\forall g \in G \cdot M_t \models goal(g,v) \tag{D.10}$$

$$\forall k, 0 \leq k \leq q \cdot g \in G, g^+ \subseteq s_k, g^- \notin s_k \Rightarrow M_t \models satisfied(g,k) \tag{D.11}$$

$$\forall n \in N \cdot M_t \models norm(n,c) \tag{D.12}$$

$$\forall n = \langle o, a_{con}, a_{sub}, dl \rangle \in N.M_t \models executed(a_{con}, t_{a_{con}}),$$
$$\forall k, t_{a_{con}} + d(a_{con}) \leq k \leq t_{a_{con}} + d(a_{con}) + dl \cdot$$
$$M_t \models holdsat(o(n, t_{a_{con}}, a_{sub}, t_{a_{con}} + d(a_{con}) + dl), k) \tag{D.13}$$

$$\exists k_2, 0 \leq k < q \cdot M_t \models holdsat(o(n, k_1, a, dl'), k_2), M_t \models executed(a, k_2),$$
$$k_2! = dl' \Rightarrow M_t \models cmp(o(n, k_1, a, dl'), k_2) \tag{D.14}$$

$$\exists k_2, 0 \leq k < q, M_t \models cmp(o(n, k_1, a, dl'), k_2) \Rightarrow$$
$$M_t \models terminated(o(n, k_1, a, dl'), k_2) \tag{D.15}$$

$$\exists k_2, k_2 = dl', M_t \models holdsat(o(n, k_1, a, dl'), k_2) \Rightarrow M_t \models vol(o(n, k_1, a, dl'), k_2) \tag{D.16}$$

$$\exists k_2, 0 \leq k_2 \leq q, M_t \models vol(o(n, k_1, a, dl'), k_2) \Rightarrow$$
$$M_t \models terminated(o(n, k_1, a, dl'), k_2) \tag{D.17}$$

$$\forall n = \langle f, a_{con}, a_{sub}, dl' \rangle \in N.M_t \models executed(a_{con}, t_{a_{con}}),$$
$$\forall k, t_{a_{con}} + d(a_{con}) \leq k \leq t_{a_{con}} + d(a_{con}) + dl' \cdot$$
$$M_t \models holdsat(f(n, t_{a_{con}}, a_{sub}, t_{a_{con}} + d(a_{con}) + dl'), k) \tag{D.18}$$

$$\exists k_2, k_2 = dl', M_t \models holdsat(f(n, k_1, a, dl'), k_2) \Rightarrow$$
$$M_t \models cmp(f(n, k_1, a, dl'), k_2) \tag{D.19}$$

$$\exists k_2, 0 \leq k_2 \leq q, M_t \models cmp(f(n, k_1, a, dl'), k_2) \Rightarrow$$
$$M_t \models terminated(f(n, k_1, a, dl'), k_2) \tag{D.20}$$

$$\forall k_2, 0 \leq k_2 < q \cdot M_t \models holdsat(f(n, k_1, a, dl'), k_2), M_t \models executed(a, k_2) \Rightarrow$$
$$M_t \models vol(f(n, k_1, a, dl'), k_2) \tag{D.21}$$

$$\exists k_2, 0 \leq k_2 \leq q, M_t \models vol(f(n, k_1, a, dl'), k_2) \Rightarrow$$
$$M_t \models terminated(f(n, k_1, a, dl'), k_2) \quad \text{(D.22)}$$

$$\exists k, 0 \leq k \leq q, M_t \models satisfied(g, k) \Rightarrow M_t \models satisfied(g) \quad \text{(D.23)}$$

We need to prove that M_t is an answer set of Π_{PBase}. Therefore, we need to demonstrate that M_t is a minimal model for $\Pi_{PBase}^{M_t}$. Let $r \in \Pi_{PBase}^{M_t}$ be an applicable rule. In order for M_t to be a model of $\Pi_{PBase}^{M_t}$, we need to show that r is applied (i.e. $M_t \models Head(r)$). We will go through each rule in the same order in Lines 1–43.

- r is of type rule in Line 1: fact and automatically applied.
- r is of type rule in Line 2: fact and automatically applied.
- r is of type rule in Lines 3–4: because of Gelfond-Lifschitz transformation, we know that *not terminated*(x, s) is removed from this rule. Combination of D.2 and D.3 for x at k gives $M_t \models holdsat(x, k+1)$.
- r is of type rule in Line 5: fact and automatically applied.
- r is of type rule in Line 6: after Gelfond-Lifschitz transformation, from the body and description of this rule we have $a \in A$ and $pr(a)^+ \in s_k$ and $pr(a)^- \notin s_k$, which with D.5 implies that $M_t \models pre(a, k)$.
- r is of type rule in Line 7: any action $a \in A$ can be executed in a state. After the transformation for choice rules, we obtain $\forall (a, k) \widehat{\in} \pi$ we have $executed(a, k) \ :- \ action(a, d), \ state(k).$ and $\forall (a, k)$ s.t. $(a, k) \ \widehat{\notin} \ \pi$ we have $:- \ action(a, d), state(k)$. With D.6, we know that $M_t \models executed(a, k)$.
- r is of type rule in Lines 8–9: *inprog* atoms originate from execution of actions. From D.6 we know that $\forall (a, k) \widehat{\in} \pi, M_t \models executed(a, k)$. Since a is executed with D.7 we have $\forall k, t_a \leq k < t_a + d(a), M_t \models inprog(a, k)$.
- r is of type rule in Line 10: the head of this rule is empty. Since π is a plan and we have assumed the preconditions of actions in a plan are hold while the actions are in progress, this rule is applied.
- r is of type rule in Lines 11–12: the head of this rule is empty. Because π is a plan and we have assumed that two actions in a plan cannot have exactly the same start state, this rule is applied.
- r is of type rule in Lines 13–14: the body of this rule implies that the add postconditions of an executed action a hold in the state in which the action ends. Since $\forall (a, k) \widehat{\in} \pi, M_t \models executed(a, k)$, with D.8 we have $\forall x \in ps(a)^+ \cdot M_t \models holdsat(x, k + d(a))$.

- r is of type rule in Lines 15–16: the body of this rule implies that the delete postconditions of an executed action a are terminated in the state before the end state of the action. Since $\forall(a,k) \,\widehat{\in}\, \pi, M_t \models executed(a,k)$, with D.9 we have $\forall x \in ps(a)^- \cdot M_t \models terminated(x, k+d(a)-1)$.
- r is of type rule in Line 17: fact and automatically applied.
- r is of type rule in Line 18: after Gelfond-Lifschitz transformation, from the body and description of this rule we have $g^+ \in s_k$ and $g^- \notin s_k$, which with D.11 implies that $M_t \models satisfied(g,k)$.
- r is of type rule in Line 19: fact and automatically applied.
- r is of type rule in Lines 20–21: the body of the rule implies that normative fluents for obligations hold over the compliance period if the action that is the condition of the norm is executed. If a_{con} for an obligation norm belongs to π, then based on D.6 we know that $M_t \models executed(a_{con}, t_{a_{con}})$. From D.6 and D.13 we know that $M_t \models holdsat(o(n,k_1,a_{sub},t_{a_{con}}+d(a_{con})+dl'),k_2)$ over the period $t_{a_{con}}+d(a_{con}) \leq k_2 \leq t_{a_{con}}+d(a_{con})+dl'$.
- r is of type rule in Lines 22–23: this rule states that if the obliged action is executed while the normative fluent holds, the norm is complied with. If a_{sub} is executed in π ($M_t \models executed(a_{con}, t_{a_{con}})$) while the normative fluent in D.13 holds, with D.14 we know that M_t models the compliance atom.
- r is of type rule in Lines 24–25: complied obligations are terminated in the compliance state. With D.14 we know that M_t models compliance atoms, and D.15 implies that they are terminated in the same state.
- r is of type rule in Lines 26–27: if the obligation fluent still holds when the deadline occurs, the obligation is violated. D.16 implies this is modelled by M_t.
- r is of type rule in Line 28–29: violated obligations are terminated in the violation state. With D.16 we know that M_t models violation atoms, and D.17 implies that they are terminated.
- r is of type rule in Lines 30–31: the body of the rule implies that normative fluents for prohibition norms are hold over the compliance period if the action that is the condition of the norm is executed. If a_{con} for a prohibition belongs to π, then based on D.6 we know that $M_t \models executed(a_{con}, t_{a_{con}})$. From D.6 and D.18 we know that $M_t \models holdsat(f(n,k_1,a_{sub},t_{a_{con}}+d(a_{con})+dl'),k_2)$ over the period $t_{a_{con}}+d(a_{con}) \leq k_2 \leq t_{a_{con}}+d(a_{con})+dl'$.

- r is of type rule in Lines 32–33: this rule states that if the normative fluent still holds at the end of compliance period, the prohibition is complied with. D.19 implies that M_t models the head of this rule.
- r is of type rule in Line 34–35: complied prohibitions are terminated in the compliance state. With D.19 we know that M_t models compliance atoms, and D.20 implies that they are terminated and this rule is applied.
- r is of type rule in Lines 36–37: this rule states that if the prohibited action is executed while the normative fluent holds, the norm is violated. If a_{sub} is executed in π ($M_t \models executed(a_{sub}, t_{a_{sub}})$) while the normative fluent in D.18 holds, with D.21 we know that M_t models the violation atom.
- r is of type rule in Line 38–39: violated prohibitions are terminated in violation state. With D.21 we know that M_t models violation atoms, and D.22 implies that they are terminated.
- r is of type rule in Line 40: this rule is applicable whenever a goal is satisfied in a state. With D.11 and D.23 we can obtain this ($M_t \models satisfied(g)$).
- r is of type rule in Line 41: the head of this rule is empty. Because π is a plan it has to satisfy at least one goal, so this rule is applied.
- r is of type rule in Line 42–43: the head of this rule is empty. Because π is a plan and a plan cannot contain concurrent execution of actions, this rule is applied.

By showing that all the rules, apart from those not applicable, are applied, we have shown that M_t is a model for $\Pi^{M_t}_{PBase}$.

Now, we need to show that M_t is minimal, which means that there exists no other model of $\Pi^{M_t}_{PBase}$ that is a subset of M_t.

Let $M \subset M_t$ be a model for $\Pi^{M_t}_{PBase}$, then there must exist an atom $s \in (M_t \setminus M)$. If s is an atom that is generated because it is a fact, it must belong to M too and if that is not the case, then M cannot be a model. We now proceed with the rest of atoms that do not appear as facts:

- $s = executed(a, k)$: $M_t \models s$ implies that $r : executed(a, k) :- action(a, d), state(k).$, $r \in \Pi^{M_t}_{PBase}$. If $M \not\models executed(a, k)$, then r was applicable but not applied, therefore, M is not a model.
- $s = holdsat(x, k)$: $M_t \models s$ implies that one of the following four condition must have occurred:
  - D.2: if this is the case then s was true in s_{k-1} (not terminated is removed from the body of this rule because of the Gelfond-

Lifschitz). In this case the construction of M_t guarantees $M_t \models holdsat(x, k-1)$. If $M \not\models holdsat(a, k)$ then rule in Lines 3–4 are applicable but not applied, so M cannot be a model.

- D.8: if this is the case then we know $M_t \models executed(a, k)$. Earlier we showed that if $M_t \models executed(a, k)$, then $M \models executed(a, k)$ too. Thus, if $M \not\models holdsat(x, k + d(a))$ for all $x \in ps(a)^+$, then rule in Lines 13–14 is applicable but not applied, so M cannot be a model.
- D.13: if this is the case then $x = o(n, k, a_{sub}, t_{a_{con}} + d(a_{con}) + dl)$. Because M_t is a model then $M_t \models executed(a_{con}, t_{a_{con}})$. So $M \models executed(a_{con}, t_{a_{con}})$ too. Therefore, rule in Lines 20–21 is applicable and if $M \not\models s$ then this rule is not applied and M cannot be a model.
- D.18: if this is the case then $x = f(n, k, a_{sub}, t_{a_{con}} + d(a_{con}) + dl)$. Similar to reasoning above but instead of rule in Lines 20–21, rule in Lines 30–31 is applicable.

- $s = pre(a, k)$: $M_t \models s$ implies that $\forall x \in pr(a)^+, x \in s_k$ and $\forall x \in pr(a)^-, x \notin s_k$. If M is a model then according to the transformed version of rule 6, $M \models pre(a, k)$. If that is not the case, then M is not a model.
- $s = inprog(a, k)$: $M_t \models s$ implies that $M_t \models executed(a, t_a)$. Therefore, $M \models executed(a, t_a)$. That means the rule in Lines 8–9 is applicable and if $\forall t_a \leq k < t_a + d(a), M \not\models inprog(a, k)$ then this applicable rule is not applied. Therefore M cannot be a model.
- $s = cmp(o(n, k_1, a, dl'), k_2)$: $M_t \models s$ implies that $M_t \models holdsat(o(n, k_1, a, dl'), k_2)$ and also $M_t \models executed(a, k_2)$. Consequently, $M \models holdsat(o(n, k_1, a, dl'), k_2)$ and $M \models executed(a, k_2)$. As a result, rule in Lines 22–23 is applicable and if $M \not\models s$, the rule is not applied and M is not a model.
- $s = vol(o(n, k_1, a, dl'), k_2)$: $M_t \models s$ implies that $M_t \models holdsat(o(n, k_1, a, dl'), k_2)$ and also $k_2 = dl'$. Because M is a model we know that $M \models holdsat(o(n, k_1, a, dl'), k_2)$. As a result, rule in Line 27 is applicable and if $M \not\models vol(o(n, k_1, a, dl'), k_2)$, the rule is not applied and M is not a model.
- $s = cmp(f(n, k_1, a, dl'), k_2)$: $M_t \models s$ implies that $M_t \models holdsat(f(n, k_1, a, dl'), k_2)$ and also $k_2 = dl'$. Because M is a model we know that $M \models holdsat(f(n, k_1, a, dl'), k_2)$ too. As a result rule in Lines 32–33 is applicable and if $M \not\models cmp(f(n, k_1, a, dl'), k_2)$, the rule is not applied

and M is not a model.

- $s = vol(f(n, k_1, a, dl'), k_2)$: $M_t \models s$ implies that $M_t \models holdsat(f(n, k_1, a, dl'), k_2)$ and also $M_t \models executed(a, k_2)$. Consequently, we know that $M \models holdsat(f(n, k_1, a, dl'), k_2)$ and $M \models executed(a, k_2)$. As a result rule in Lines 36–37 is applicable and if $M \not\models vol(f(n, k_1, a, dl'), k_2)$, the rule is not applied and M is not a model.
- $s = terminated(x, k)$: $M_t \models s$ implies that one of the following five situations hold:
  - D.9: if this is the case then $M_t \models executed(a, k - d(a) + 1)$. If M is a model $M \models executed(a, k - d(a) + 1)$. Thus, the rule in Lines 15–16 is applicable and if $M \not\models terminated(s, k)$ then the rule is applicable but not applied, which means that M is not a model.
  - D.15 and D.17: if this is the case $x = o(n, k_2, a, dl') \in s_k$, and $M_t \models cmp(o(n, k_1, a, dl'), k)$ or $M_t \models vol(o(n, k_1, a, dl'), k)$. Since M is a model, then if $M_t \models cmp(o(n, k_1, a, dl'), k)$ the same applies to M and if $M_t \models vol(o(n, k_1, a, dl'), k)$ again the same applies to M. In either case, according to rule in Lines 24–25 and rule in Lines 28–29, $M \models terminated(o(n, k_1, a, dl'), k)$ or M is not a model.
  - D.20 and D.22: if this is the case $x = f(n, k_1, a, dl') \in s_k$, and $M_t \models cmp(f(n, k_1, a, dl'), k)$ or $M_t \models vol(f(n, k_1, a, dl'), k)$. Since M is a model, if $M_t \models cmp(f(n, k_1, a, dl'), k)$ the same applies to M and if $M_t \models vol(f(n, k_1, a, dl'), k)$ again the same applies to M. In either case, according to rule in Lines 34–35 and rule in Lines 38–39, $M \models terminated(f(n, k_1, a, dl'), k)$ or M is not a model.
- $s = satisfied(g, k)$: $M_t \models s$ implies that $\forall x \in g^+, x \in s_k$ and $\forall x \in g^-, x \notin s_k$. If M is a model then according to the transformed version of rule 18 $M \models satisfied(g, k)$ and if that is not the case then M is not a model.
- $s = satisfied(g)$: $M_t \models s$ because $M_t \models satisfied(g, k)$. Since $M_t \models satisfied(g, k)$, according to previous item $M \models satisfied(g, k)$ too, therefore rule 40 is applicable and $M \models satisfied(g)$ or it is not a model.

The combination of these items demonstrates that M cannot be a model for $\Pi^{M_t}_{PBase}$ if it differs from M_t. M_t is therefore a minimal model for $\Pi^{M_t}_{PBase}$ and an answer set for Π_{PBase}. □

References

Ågotnes, T., van der Hoek, W., Wooldridge, M., 2007. Normative system games. In: Durfee, E. H., Yokoo, M., Huhns, M. N., Shehory, O. (Eds.), International Conference on Autonomous Agents and Multiagent Systems. IFAAMAS, p. 129.

Aldewereld, H., Dignum, F., García-Camino, A., Noriega, P., Rodríguez-Aguilar, J. A., Sierra, C., 2006. Operationalisation of norms for usage in electronic institutions. In: Nakashima, H., Wellman, M. P., Weiss, G., Stone, P. (Eds.), International Conference on Autonomous Agents and Multiagent Systems. ACM, pp. 223–225.

Alechina, N., Dastani, M., Logan, B., 2012. Programming norm-aware agents. In: van der Hoek, W., Padgham, L., Conitzer, V., Winikoff, M. (Eds.), International Conference on Autonomous Agents and Multiagent Systems. IFAAMAS, pp. 1057–1064.

Alrawagfeh, W., Meneguzzi, F., 2014. Utilizing permission norms in BDI practical normative reasoning. In: International Workshop on Coordination, Organizations, Institutions, and Norms. Springer International Publishing, pp. 1–18.

Artikis, A., Sergot, M. J., Pitt, J. V., 2009. Specifying norm-governed computational societies. ACM Transactions on Computational Logic 10 (1), 1–42.

Baral, C., 2003. Knowledge Representation, Reasoning, and Declarative Problem Solving. Cambridge University Press, New York, NY, USA.

Blum, A. L., Furst, M. L., 1997. Fast planning through planning graph analysis. Artificial Intelligence 90 (1), 281 – 300.

Bordini, R. H., Wooldridge, M., Hübner, J. F., 2007. Programming Multi-Agent Systems in AgentSpeak using Jason. Wiley Series in Agent Technology. John Wiley & Sons.

Börger, E., Stärk, R., 2003. Asynchronous multi-agent asms. In: Abstract State Machines. Springer Berlin Heidelberg, pp. 207–282.

Broersen, J., Dastani, M., Hulstijn, J., Huang, Z., van der Torre, L., 2001. The BOID architecture: Conflicts between beliefs, obligations, intentions and desires. International Conference on Autonomous Agents. ACM, pp. 9–16.

Broersen, J., Dastani, M., Hulstijn, J., van der Torre, L., 2002. Goal generation in the BOID architecture. Cognitive Science Quarterly 2 (3-4), 428–447.

Chesani, F., Mello, P., Montali, M., Torroni, P., 2013. Representing and monitoring social commitments using the event calculus. Autonomous Agents and Multi-Agent Systems 27 (1), 85–130.

Cliffe, O., De Vos, M., Padget, J. A., 2006. Answer set programming for representing and reasoning about virtual institutions. In: Inoue, K., Satoh, K., Toni, F. (Eds.), International Workshop on Computational Logic in Multi-Agent Systems. Springer, pp. 60–79.

Criado, N., 2012. Using norms to control open multi-agent systems. Ph.D. thesis, Universidad Politecnica de Valencia.

Criado, N., Argente, E., Botti, V. J., 2010. A BDI architecture for normative decision making. In: van der Hoek, W., Kaminka, G. A., Lespérance, Y., Luck, M., Sen, S. (Eds.), International Conference on Autonomous Agents and Multiagent Systems. IFAAMAS, pp. 1383–1384.

Dastani, M., 2008. 2APL: a practical agent programming language. Autonomous Agents and Multi-Agent Systems 16(3), 214–248.

de Boer, F. S., Hindriks, K. V., van der Hoek, W., Meyer, J. C., 2002. Agent programming with declarative goals. CoRR cs.AI/0207008.

De Vos, M., Balke, T., Satoh, K., 2013. Combining event-and state-based norms. In: International conference on Autonomous Agents and Multi-Agent Systems. pp. 1157–1158.

Doherty, P., Gustafsson, J., Karlsson, L., Kvarnström, J., 1998. TAL: temporal action logics language specification and tutorial. Electronic Transactions on Artificial Intelligence (2), 273–306.

Eiter, T., Faber, W., Leone, N., Pfeifer, G., 1999. The diagnosis frontend of the DLV system. AI Communications 12, 12–1.

Esteva, M., Rodríguez-Aguilar, J. A., Sierra, C., Garcia, P., Arcos, J. L., 2001. On the formal specifications of electronic institutions. In: Dignum, F., Sierra, C. (Eds.), Agent Mediated Electronic Commerce, The European AgentLink Perspective. Vol. 1991 of Lecture Notes in Computer Science. Springer, pp. 126–147.

Fikes, R. E., Nilsson, N. J., 1971. Strips: A new approach to the application of theorem proving to problem solving. International Joint Conference on Artificial Intelligence. Morgan Kaufmann Publishers Inc., pp. 608–620.

Fox, M., Long, D., 2003. PDDL2.1: An extension to pddl for expressing temporal planning domains. Artificial Intelligence Research 20, 61–124.

Garrido, A., Fox, M., Long, D., 2002. A temporal planning system for durative actions of PDDL2.1. In: van Harmelen (2002), pp. 586–590.

Gasparini, L., Norman, T. J., Kollingbaum, M. J., Chen, L., Meyer, J. C., 2015. Verifying normative system specification containing collective imperatives and deadlines. In: Weiss, G., Yolum, P., Bordini, R. H., Elkind, E. (Eds.), International Conference on Autonomous Agents and Multiagent Systems. ACM, pp. 1821–1822.

Gebser, M., Kaminski, R., Kaufmann, B., Ostrowski, M., Schaub, T., Schneider, M., 2011. Potassco: The Potsdam answer set solving collection. AI Communications 24 (2), 107–124.

Gelfond, M., Lifschitz, V., 1988. The stable model semantics for logic programming. In: Kowalski, R. A., Bowen, K. A. (Eds.), International Conference and Symposium on Logic Programming. MIT Press, pp. 1070–1080.

Gelfond, M., Lifschitz, V., 1998. Action languages. Electronic Transactions on Artificial Intelligence 2, 93–210.

van Harmelen, F. (Ed.), 2002. Proceedings of the 15th Eureopean Conference on Artificial Intelligence, ECAI'2002, Lyon, France, July 2002. IOS Press.

Hoffmann, J., Nebel, B., 2001. The FF planning system: Fast plan generation through heuristic search. Artificial Intelligence Research, 14, 253–302.

Hulstijn, J., van der Torre, L. W. N., 2004. Combining goal generation and planning in an argumentation framework. In: Delgrande, J. P., Schaub, T. (Eds.), Non Monotonic Reasoning. pp. 212–218.

Kafali, Ö., Günay, A., Yolum, P., 2014. GOSU: computing goal support with commitments in multiagent systems. In: Schaub, T., Friedrich, G., O'Sullivan, B. (Eds.), European Conference on Artificial Intelligence. Vol. 263 of Frontiers in Artificial Intelligence and Applications. IOS Press, pp. 477–482.

Knoblock, C. A., 1994. Generating parallel execution plans with a partial-order planner. In: Hammond, K. J. (Ed.), International Conference on Artificial Intelligence Planning Systems. AAAI, pp. 98–103.

Kollingbaum, M., 2005. Norm-governed practical reasonig agents. Ph.D. thesis, University of Aberdeen.

Kowalski, R., Sergot, M., Jan. 1986. A logic-based calculus of events. New Generation Computing. 4(1), 67–95.

Lee, J., Lifschitz, V., Palla, R., 2008. A reductive semantics for counting and choice in answer set programming. In: Fox, D., Gomes, C. P. (Eds.), International Conference on Artificial Intelligence. AAAI Press, pp.472–479.

Lee, J., Padget, J., Logan, B., Dybalova, D., Alechina, N., 2014. N-Jason: Run-time norm compliance in AgentSpeak(L). In: Dalpiaz, F., Dix, J., van Riemsdijk, M. B. (Eds.), International Workshop in Engineering Multi-Agent Systems. Vol. 8758 of Lecture Notes in Computer Science. Springer, pp. 367–387.

Lee, J., Palla, R., 2012. Reformulating temporal action logics in answer set programming. In: Hoffmann, J., Selman, B. (Eds.), Conference on Artificial Intelligence. AAAI Press, 786–792.

Lee, J., Palla, R., 2014. Reformulating the situation calculus and the event calculus in the general theory of stable models and in answer set programming. CoRR abs/1401.4607.

Lifschitz, V., 2002. Answer set programming and plan generation. Artificial Intelligence 138(1-2), 39–54.

Lifschitz, V., 2008. What is answer set programming? In: Fox, D., Gomes, C. P. (Eds.), International Conference on Artificial Intelligence. AAAI, pp. 1594–1597.

López y López, F., Luck, M., d'Inverno, M., 2005. A normative framework for agent-based systems. In: Normative Multi-Agent Systems (NORMAS). pp. 24–35.

Law, M., Russo, A., Broda, k., 2015. Simplified reduct for choice rules in ASP . Technical Report., Imperial College.

Meneguzzi, F., Rodrigues, O., Oren, N., Vasconcelos, W. W., Luck, M., 2015. BDI reasoning with normative considerations. Engineering Application of Artificial Intelligence 43, 127–146.

Modgil, S., Luck, M., 2008. Argumentation based resolution of conflicts between desires and normative goals. In: Rahwan, I., Moraitis, P. (Eds.), Argumentation in Multi-Agent Systems. Vol. 5384 of Lecture Notes in Computer Science. Springer, pp. 19–36.

Nigam, V., Leite, J., 2006. A dynamic logic programming based system for agents with declarative goals. In: Baldoni, M., Endriss, U. (Eds.), International Workshop on Declarative Agent Languages and Technologies. Vol. 4327 of Lecture Notes in Computer Science. Springer, pp. 174–190.

Nunes, I., Fiadeiro, J. L., Turski, W. M., 1997. Coordination durative actions. In: Garlan, D., Métayer, D. L. (Eds.), International Conference on Coordination Languages and Models. Vol. 1282 of Lecture Notes in Computer Science. Springer, pp. 115–130.

Oren, N., Croitoru, M., Miles, S., Luck, M., 2010. Understanding permissions through graphical norms. In: Omicini, A., Sardiña, S., Vasconcelos, W. W. (Eds.), Declarative Agent Languages and Technologies. Vol. 6619 of Lecture Notes in Computer Science. Springer, pp. 167–184.

Oren, N., Vasconcelos, W., Meneguzzi, F., Luck, M., 2011. Acting on norm constrained plans. In: Leite, J., Torroni, P., Ågotnes, T., Boella, G., van der Torre, L. (Eds.), Computational Logic in Multi-Agent Systems. Vol. 6814 of Lecture Notes in Computer Science. Springer, pp. 347–363.

Panagiotidi, S., Vázquez-Salceda, J., Dignum, F., 2012a. Reasoning over norm compliance via planning. In: Aldewereld, H., Sichman, J. S. (Eds.), International Workshop on Coordination, Organizations, Institutions, and Norms in Agent Systems. Vol. 7756 of Lecture Notes in Computer Science. Springer, pp. 35–52.

Panagiotidi, S., Vázquez-Salceda, J., Vasconcelos, W., 2012b. Contextual norm-based plan evaluation via answer set programming. In: Pérez, J. B., Sánchez, M. A., Mathieu, P., Rodríguez, J. M. C., Adam, E., Ortega, A., García, M. N. M., Navarro, E., Hirsch, B., Cardoso, H. L., Julián, V. (Eds.), International Conference on Practical Applications of Agents and Multi-Agent Systems. Vol. 156 of Advances in Soft Computing. Springer, pp. 197–206.

Pitt, J., Busquets, D., Riveret, R., Sept 2013. Formal models of social processes: The pursuit of computational justice in self-organising multi-agent systems. In: International Conference on Self-Adaptive and Self-Organizing Systems. pp. 269–270.

Pokahr, A., Braubach, L., Lamersdorf, W., 2005. A goal deliberation strategy for BDI agent systems. In: Eymann, T., Klügl, F., Lamersdorf, W., Klusch, M., Huhns, M. N. (Eds.), German Conference on Multiagent System Technologies. Vol. 3550 of Lecture Notes in Computer Science. Springer, pp. 82–93.

Rahwan, I., Amgoud, L., 2006. An argumentation-based approach for practical reasoning. In: Maudet, N., Parsons, S., Rahwan, I. (Eds.), International Workshop on Argumentation in Multi-Agent Systems. Vol. 4766 of Lecture Notes in Computer Science. Springer, pp. 74–90.

Rao, A. S., Georgeff, M. P., 1995. BDI agents: From theory to practice. In: International Conference On Multi-Agent Systems. pp. 312–319.

van Riemsdijk, M. B., van der Hoek, W., Meyer, J. C., 2002. Agent programming in dribble: From beliefs to goals with plans. In: Hinchey, M. G., Rash, J. L., Truszkowski, W., Rouff, C., Gordon-Spears, D. F. (Eds.), Formal Approaches to Agent-Based Systems. Vol. 2699 of Lecture Notes in Computer Science. Springer, pp. 294–295.

van Riemsdijk, M. B., Dastani, M., Winikoff, M., 2008. Goals in agent systems: a unifying framework. In: Padgham, L., Parkes, D. C., Müller, J. P., Parsons, S. (Eds.), 7th International Joint Conference on Autonomous Agents and Multiagent Systems, 2008, Volume 2. IFAAMAS, pp. 713–720.

van Riemsdijk, M. B., Dastani, M., Meyer, J. C., 2009. Goals in conflict: semantic foundations of goals in agent programming. Autonomous Agents and Multi-Agent Systems 18 (3), 471–500.

Sartor, G., 1992. Normative conflicts in legal reasoning. Artificial Intelligence and Law 1 (2-3), 209–235.

Savarimuthu, B. T. R., Padget, J., Purvis, M. A., 2013. Social norm recommendation for virtual agent societies. In: International Conference on Principles and Practice of Multi-Agent Systems. Springer, pp. 308–323.

Shams, Z., 2016. An Argumentation-Based Approach to Normative Practical Reasoning. Ph.D. thesis, University of Bath.

Shams, Z., De Vos, M., Padget, J., Vasconcelos, W. 2015. Implementation of normative practical reasoning with durative actions. In: Coordination, Organizations, Institutions, and Norms in Agent Systems. Springer, pp. 136–154.

Thangarajah, J., Padgham, L., Winikoff, M., 2003. Detecting & avoiding interference between goals in intelligent agents. In: Gottlob, G., Walsh, T. (Eds.), Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, 2003. Morgan Kaufmann, pp. 721–726.

Thangarajah, J., Winikoff, M., Padgham, L., Fischer, K., 2002. Avoiding resource conflicts in intelligent agents. In: van Harmelen (2002), pp. 18–22.

To, S. T., Son, T. C., Pontelli, E., 2015. A generic approach to planning in the presence of incomplete information: Theory and implementation. Artificial Intelligence 227, 1–51.

Toniolo, A., 2013. Models of argument for deliberative dialogue in complex domains. Ph.D. thesis, University of Aberdeen.

Vasconcelos, W. W., Kollingbaum, M. J., Norman, T. J., 2009. Normative conflict resolution in multi-agent systems. Autonomous Agents and Multi-Agent Systems 19 (2), 124–152.